

IPv6 Motivations and Obstacles

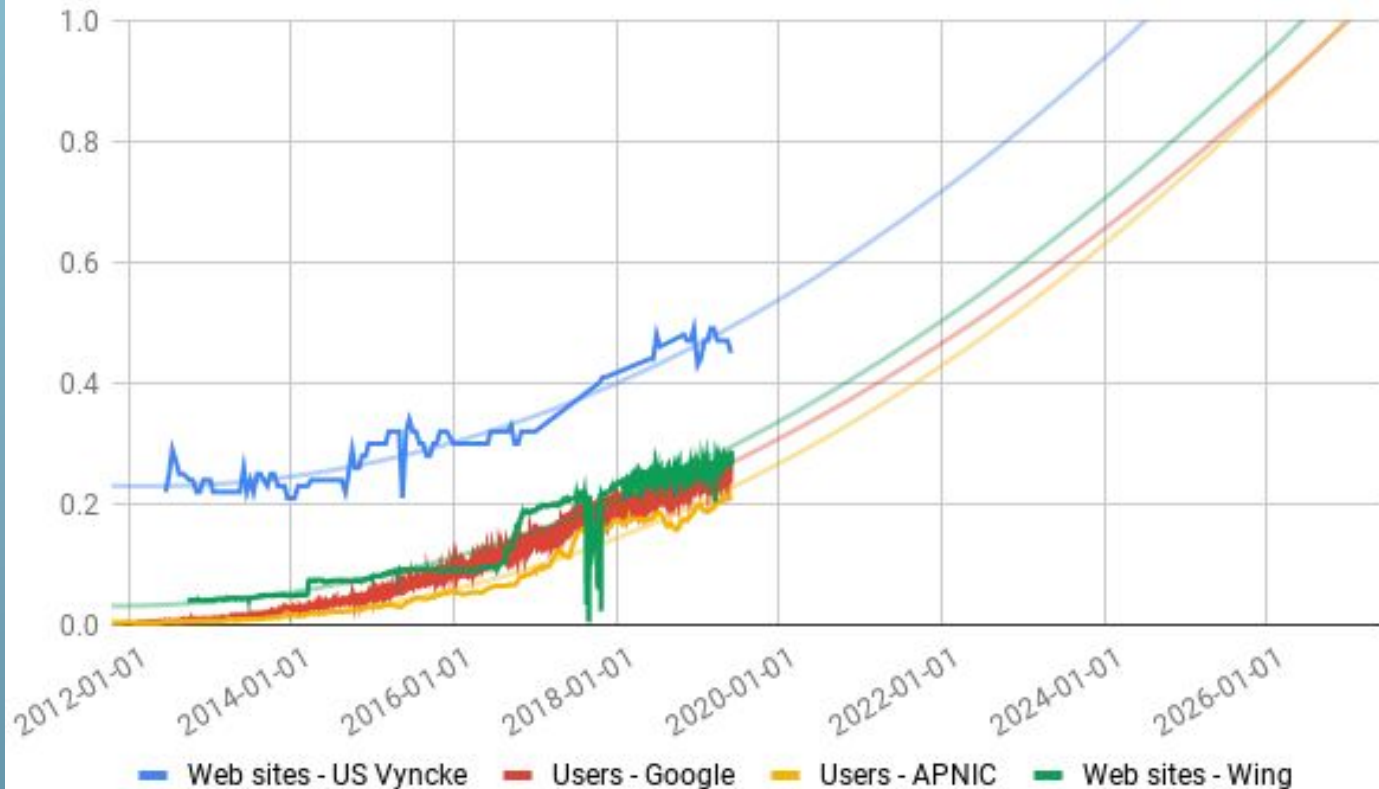
Lee Howard



Agenda

- Economic Drivers
- Cool new technologies
- External Factors
- Perceived obstacles
- Top obstacles
- IPv6-only

IPv6 Growth



Economic Drivers

IPv6 Speed

APNIC 2013

Cisco 2014

TWC 2014

Akamai 2016

LinkedIn 2016

Facebook 2017

Bajpai, Schönwälder 2017

APNIC yesterday

IPv6 is faster more often than IPv4 is.

IPv6 is faster more often than IPv4 is.

IPv6 is 10% faster on average.

(iPhone/VzW) 95% sites are 15% faster.

IPv6 is often 15-25% faster.

IPv6 is 30-40% (or less) faster.

95% of sites are same or faster.

In most regions, IPv6 is 20ms faster.

Value of a Millisecond

“Every 100ms of latency costs 1% in Sales”

Amazon

“100-millisecond delay in website load time can hurt conversion rates by 7%”

Google

Akamai

“Traffic and revenue ... dropped by 20%. . . Half a second delay caused a 20% drop in traffic.”

Value of a Millisecond

20ms =
+ 0.2% in sales
= \$400 million

Amazon

20ms =
+ 1.4% in sales =
\$38 million

Google

Akamai

$\frac{1}{2}$ sec = 20% in revenue
= \$1.1 billion

Cost of CGN

Hardware:
\$1000/Gbps

400 users

\$3800

**\$9.50 per
user**

Systems updates:
\$800

IPv4 Addresses:
\$2000/Gbps

Technology Drivers

PDM

- Sender includes in a DestinationOptions Header:
 - Packet Sequence # this packet
 - Packet Sequence # last received
 - Time between last packet sent and last received
 - Time between last packet received and last sent
- Allows you to determine RTT and server delay

rfc8250 "IPv6 Performance and Diagnostic Metrics (PDM) Destination Option"
Ackermann et al.

M-PDM

- If implemented, will provide:
 - Delay generated by this host
 - Delay generated by remote host
 - Sequence numbers for reading in packet captures
- Proposed HBH option will let middleboxes add their own information

<https://tools.ietf.org/html/draft-fear-ippm-mpdm-01> combines rfc8250 “IPv6 Performance and Diagnostic Metrics (PDM) Destination Option” and rfc8321 “Alternate-Marking Method for Passive and Hybrid Performance Monitoring”

Reserving Bits

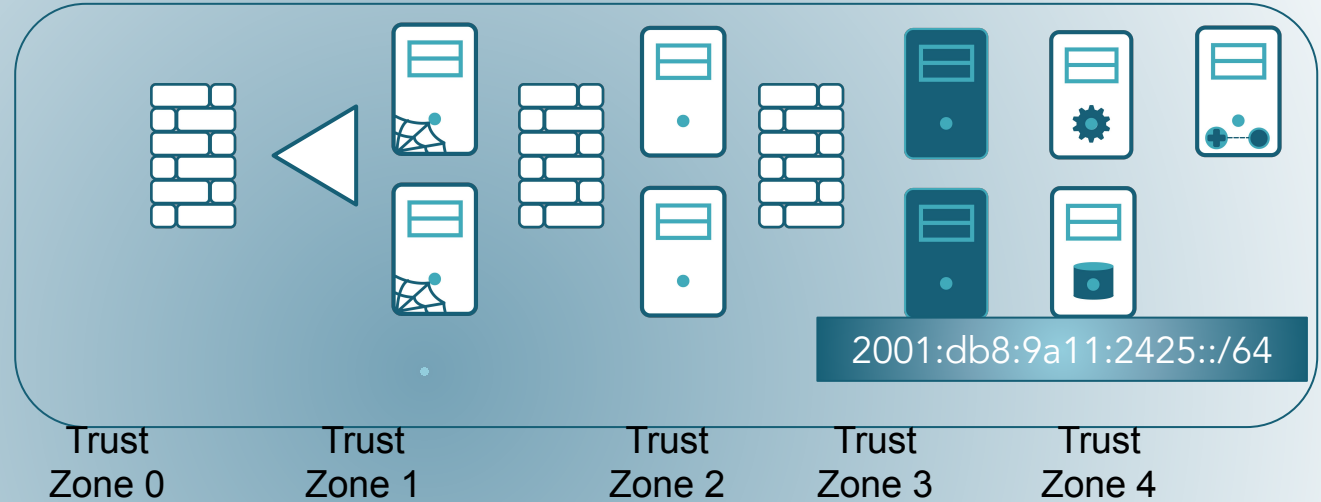
2001:db8:xxRR::/48
2001:db8:xxRR:DTAA::/64

R = Region 0-255

D = Data Center 0-15

T = Trust Zone 0-15

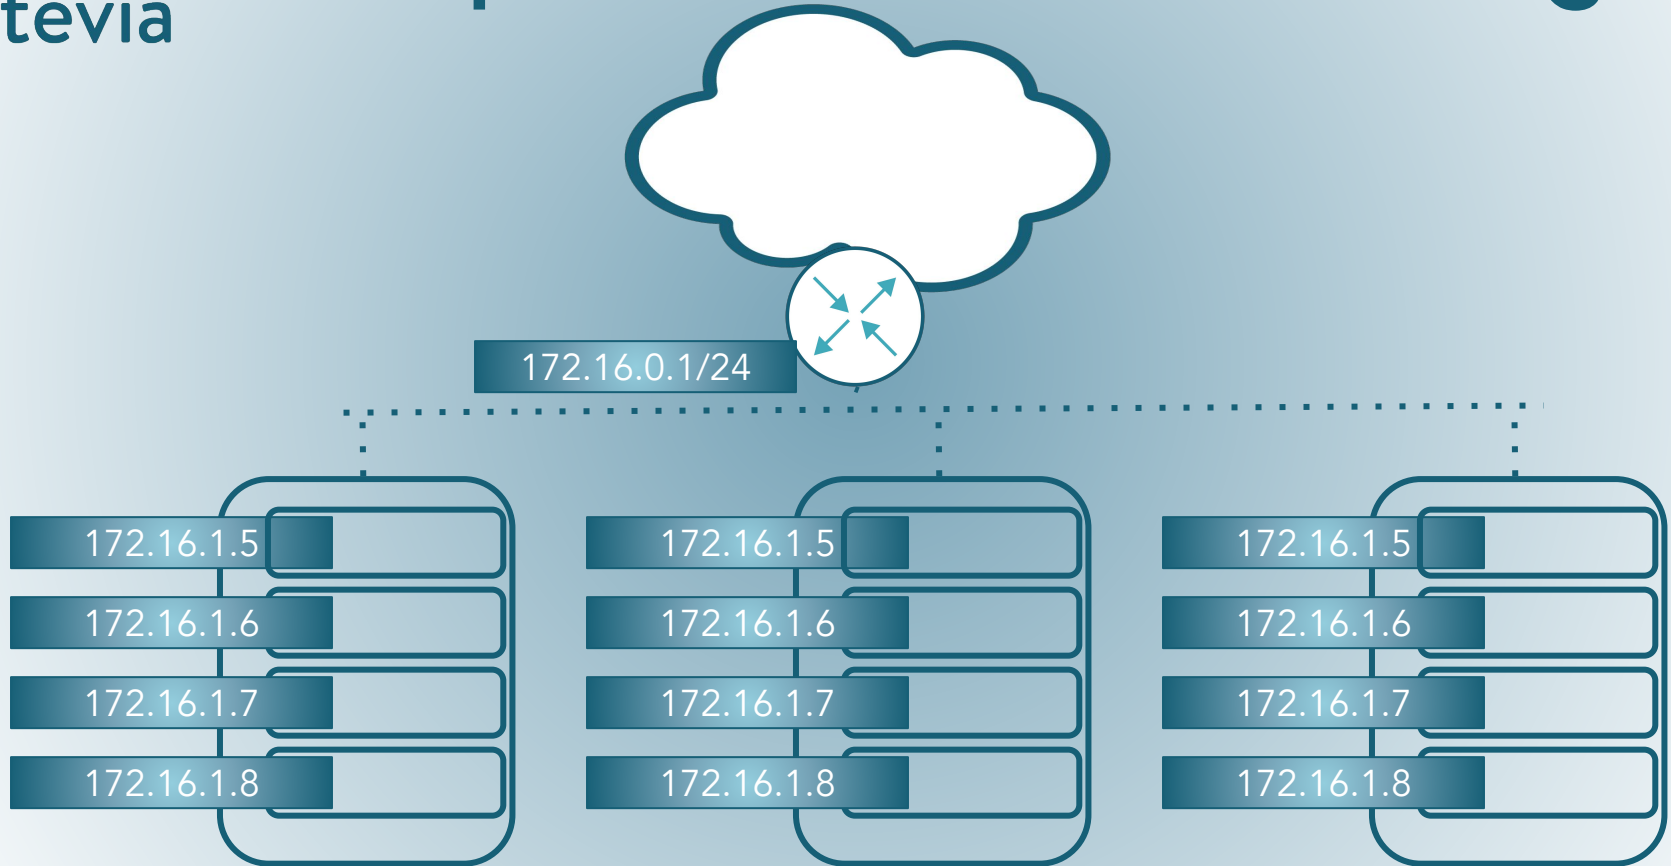
AA = Application 0-255



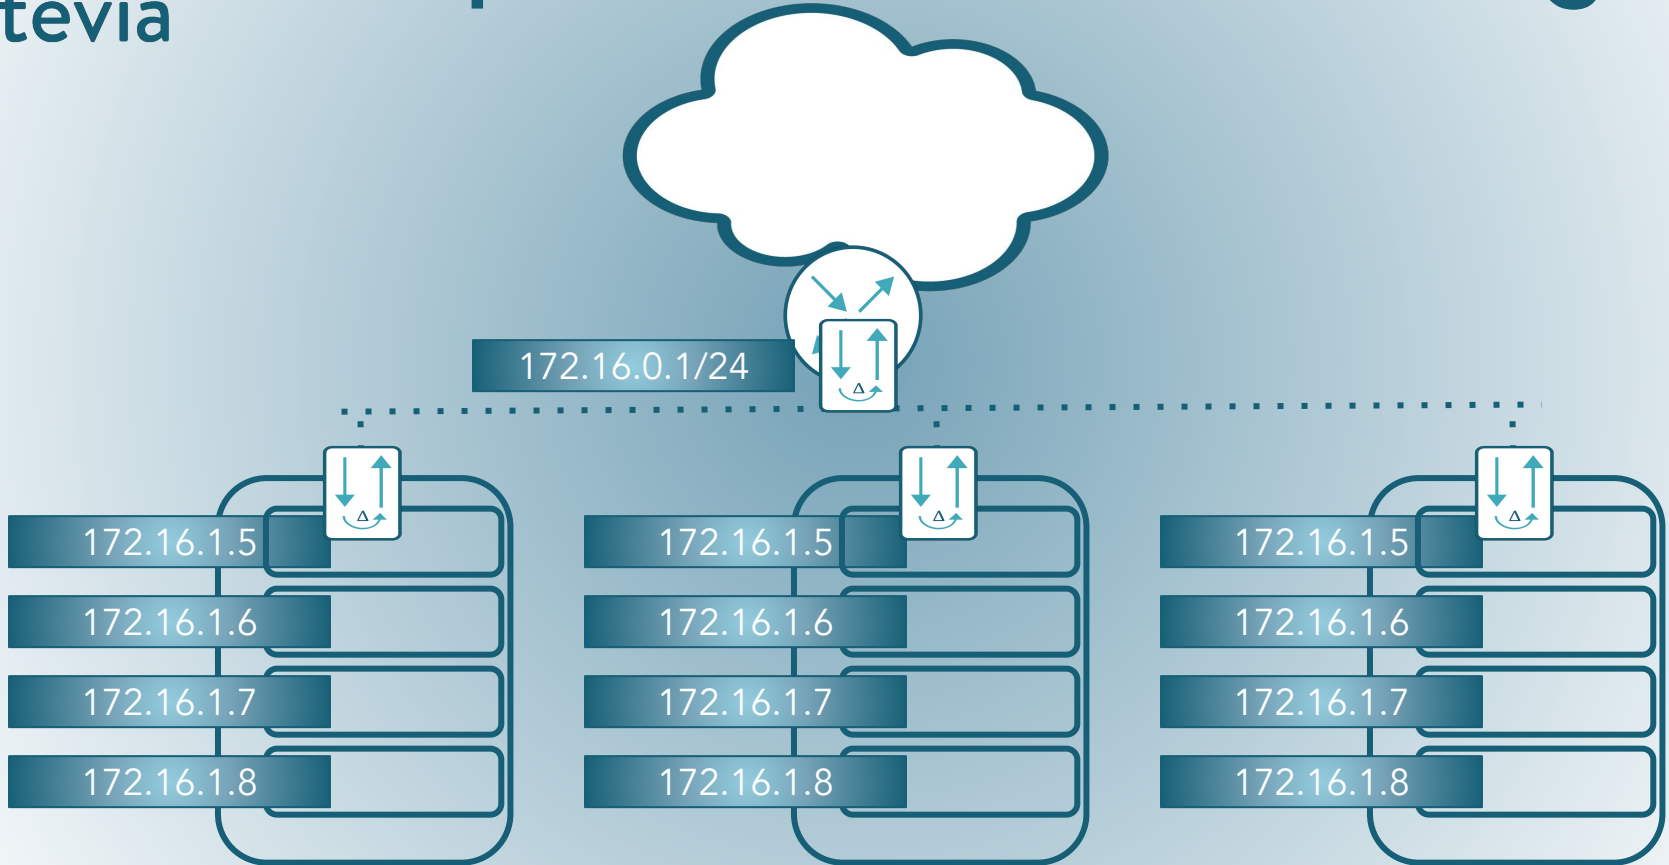
Region 17 (0x11), Data Center 2,
Trust Zone 4, Application 25

2001:db8:9a11:2425:0123:4567:89ab:cdef₁₂

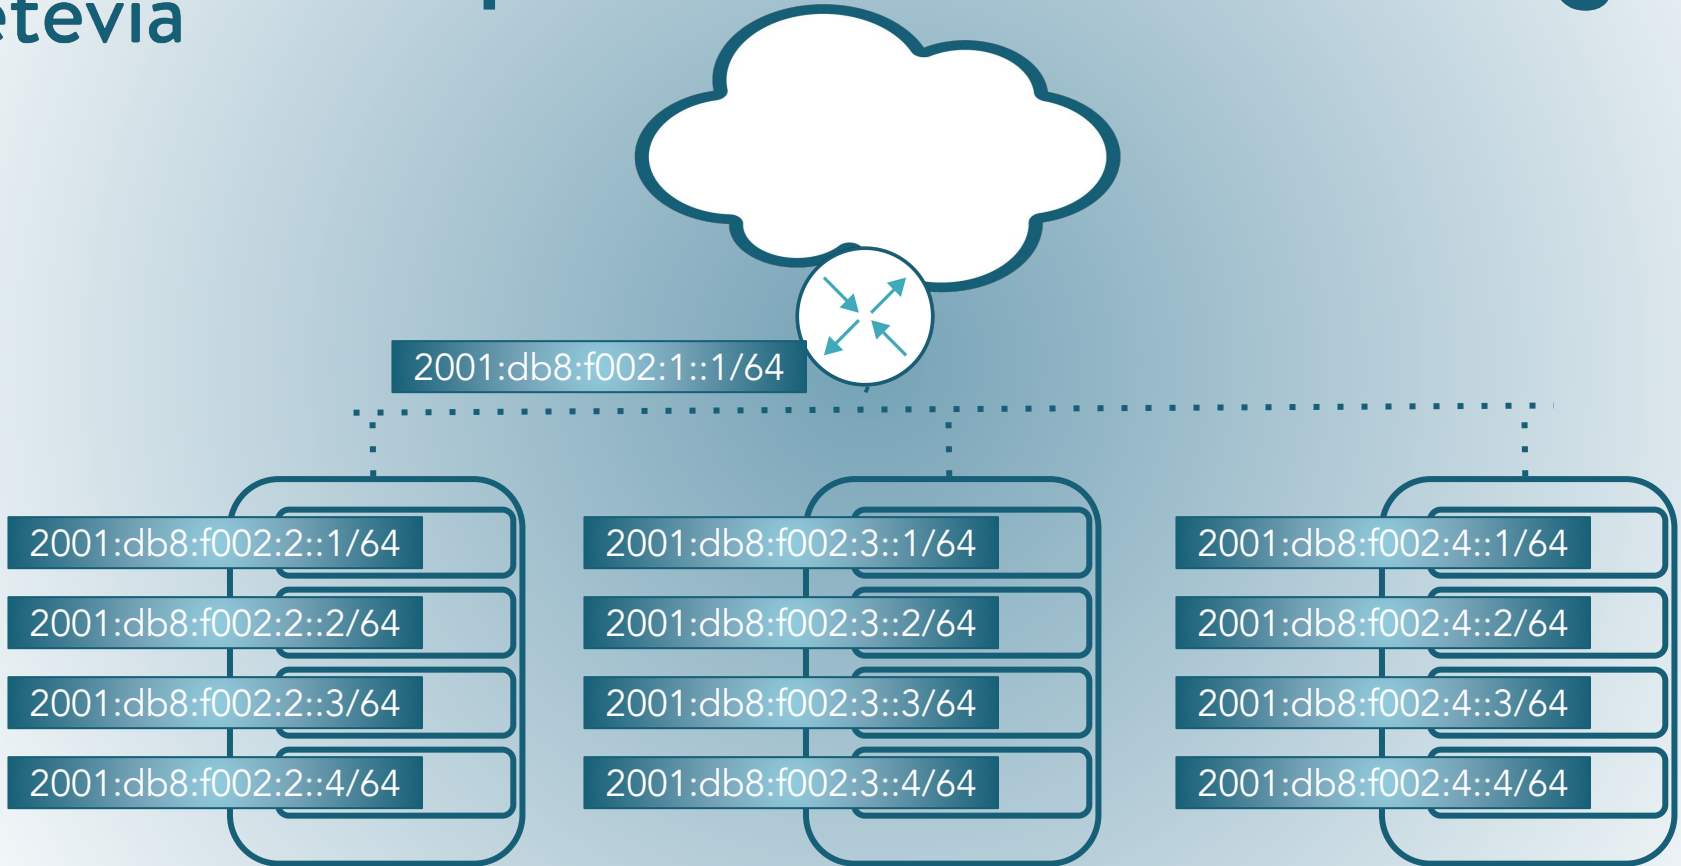
Simpler Container Numbering



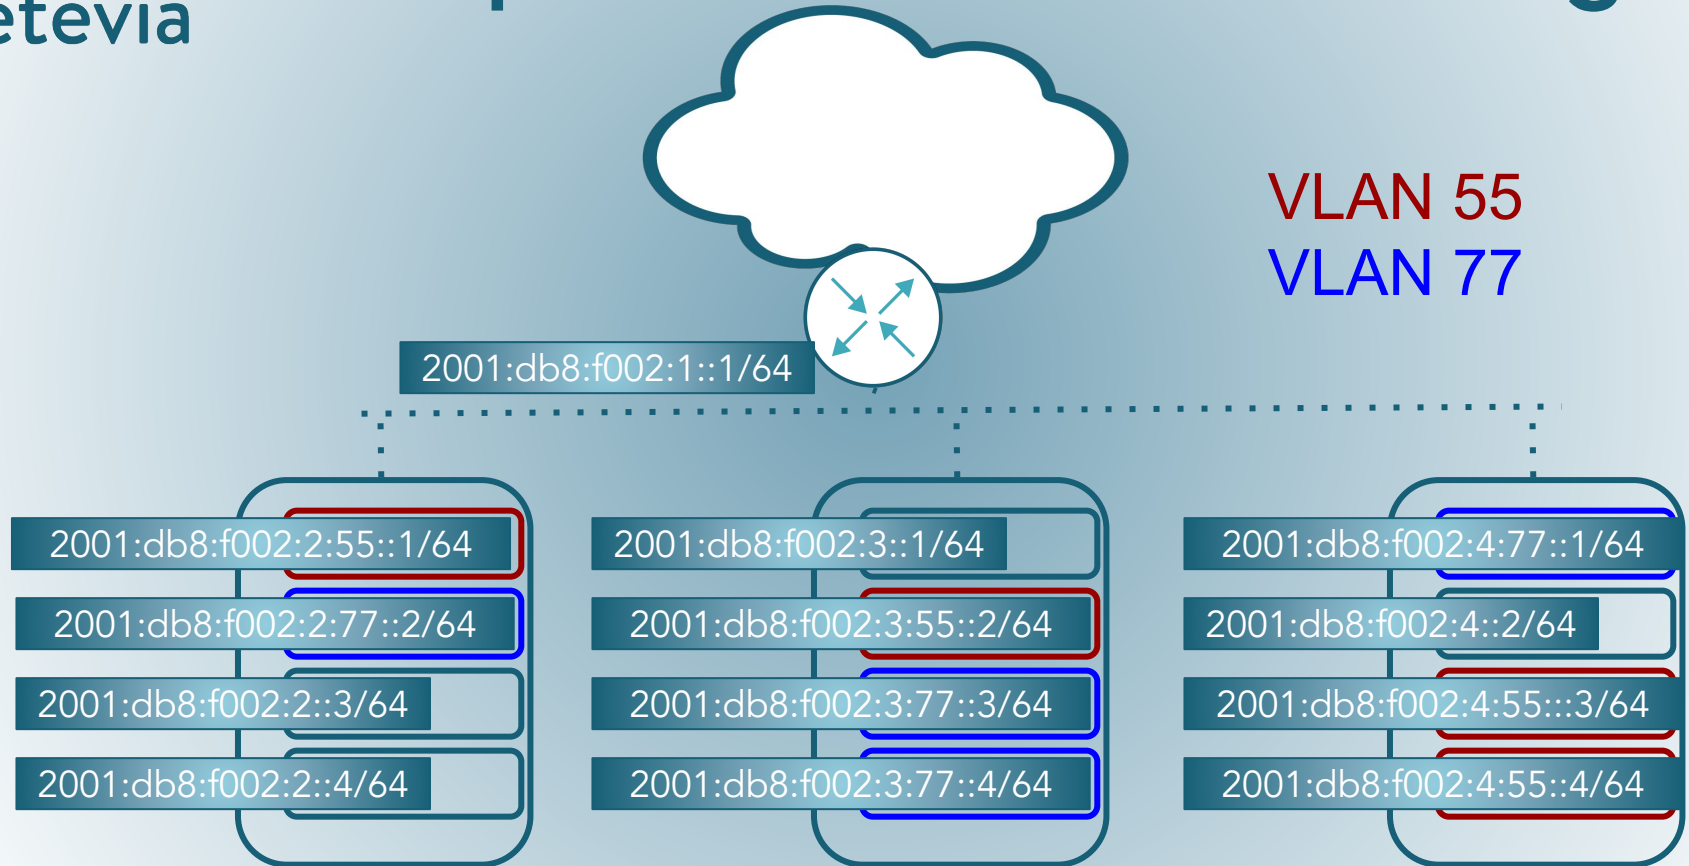
Simpler Container Numbering



Simpler Container Numbering



Simpler Container Numbering



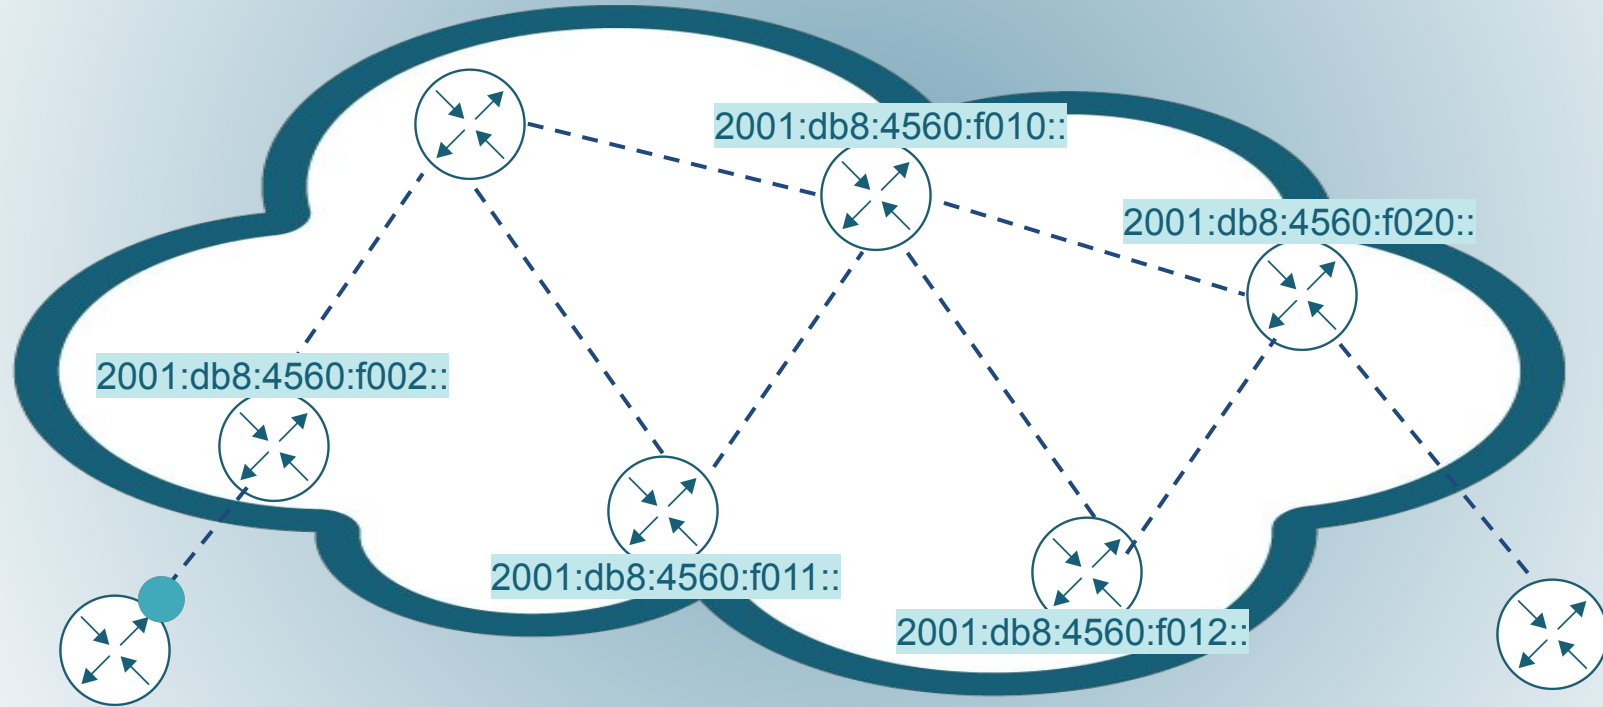
VLAN 55
VLAN 77

Finer Control over Routing

Sometimes

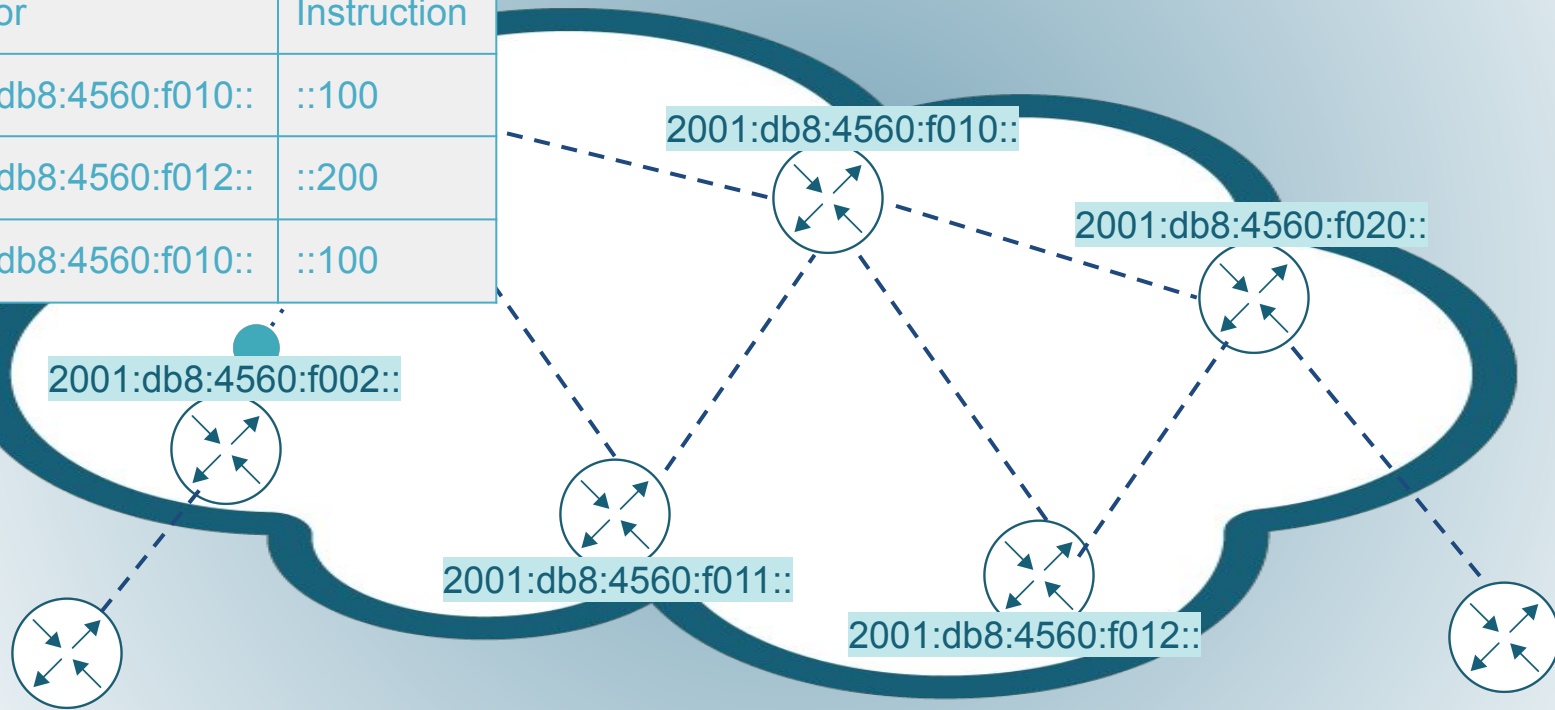
- “best path” != “shortest path”
- You want to abstract the path
- You want to avoid per-flow state
- You want to have a backup route pre-calculated (FRR)
- Lots of protocols make things complicated
- You want NFV

Segment Routing (SRv6)

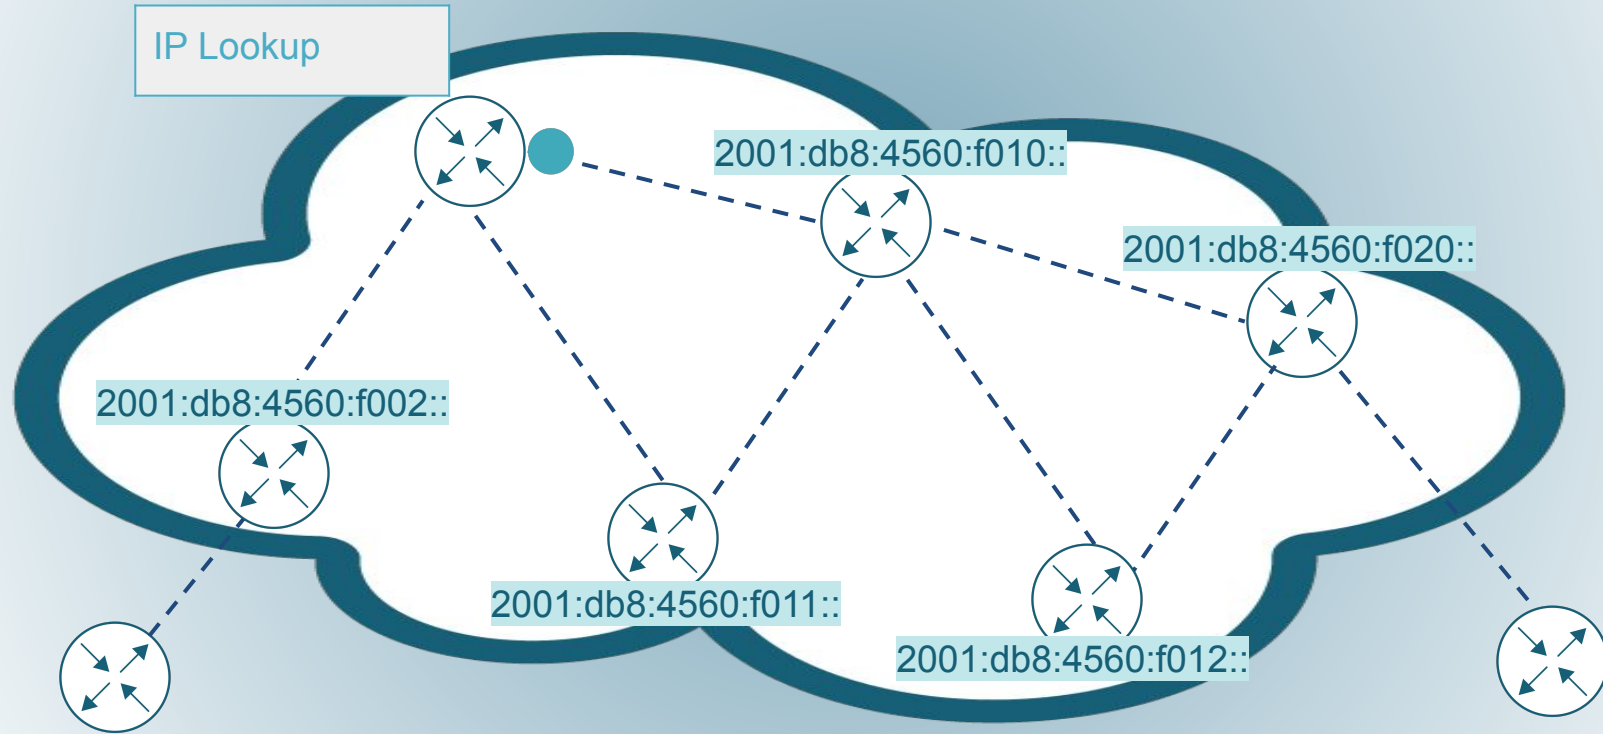


Segment Routing (SRv6)

Locator	Instruction
2001:db8:4560:f010::	::100
2001:db8:4560:f012::	::200
2001:db8:4560:f010::	::100

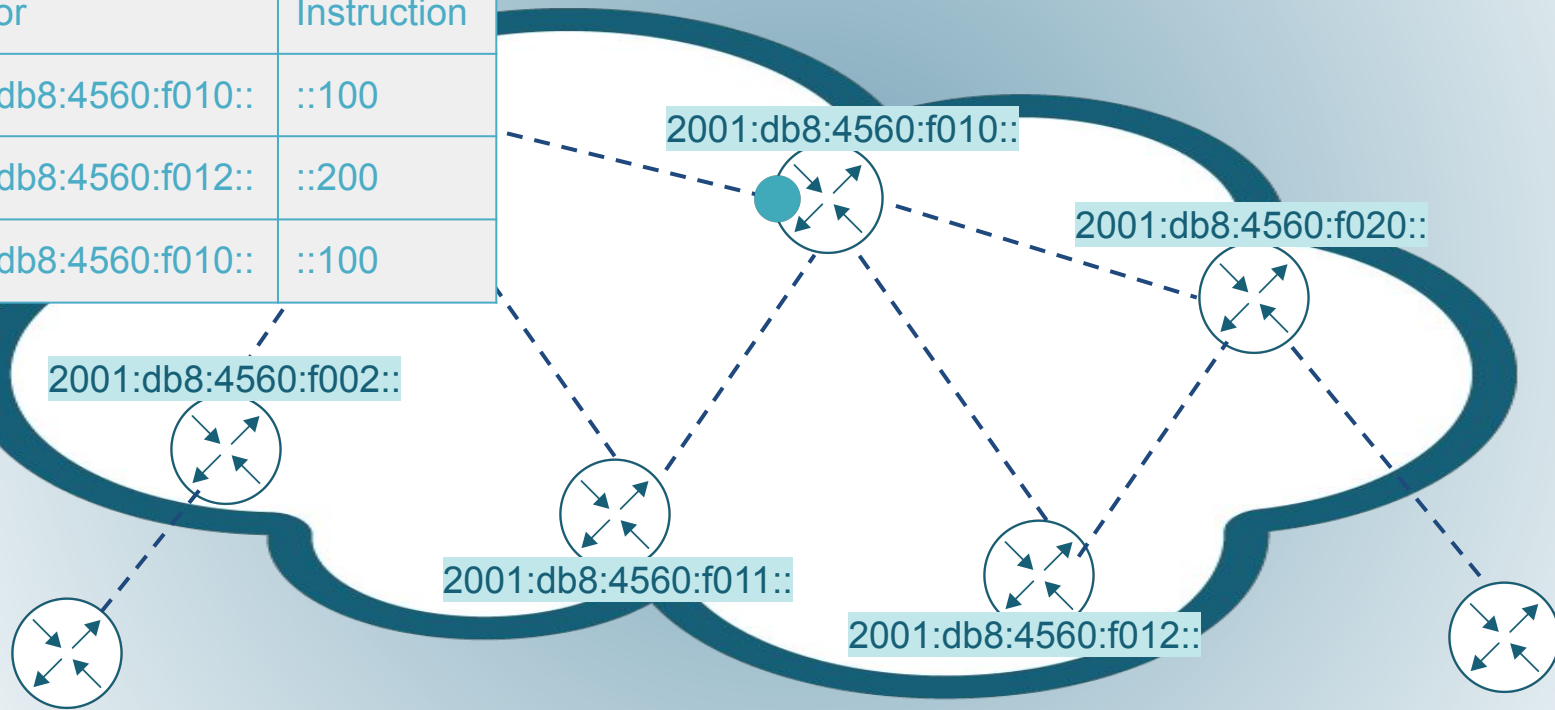


Segment Routing (SRv6)



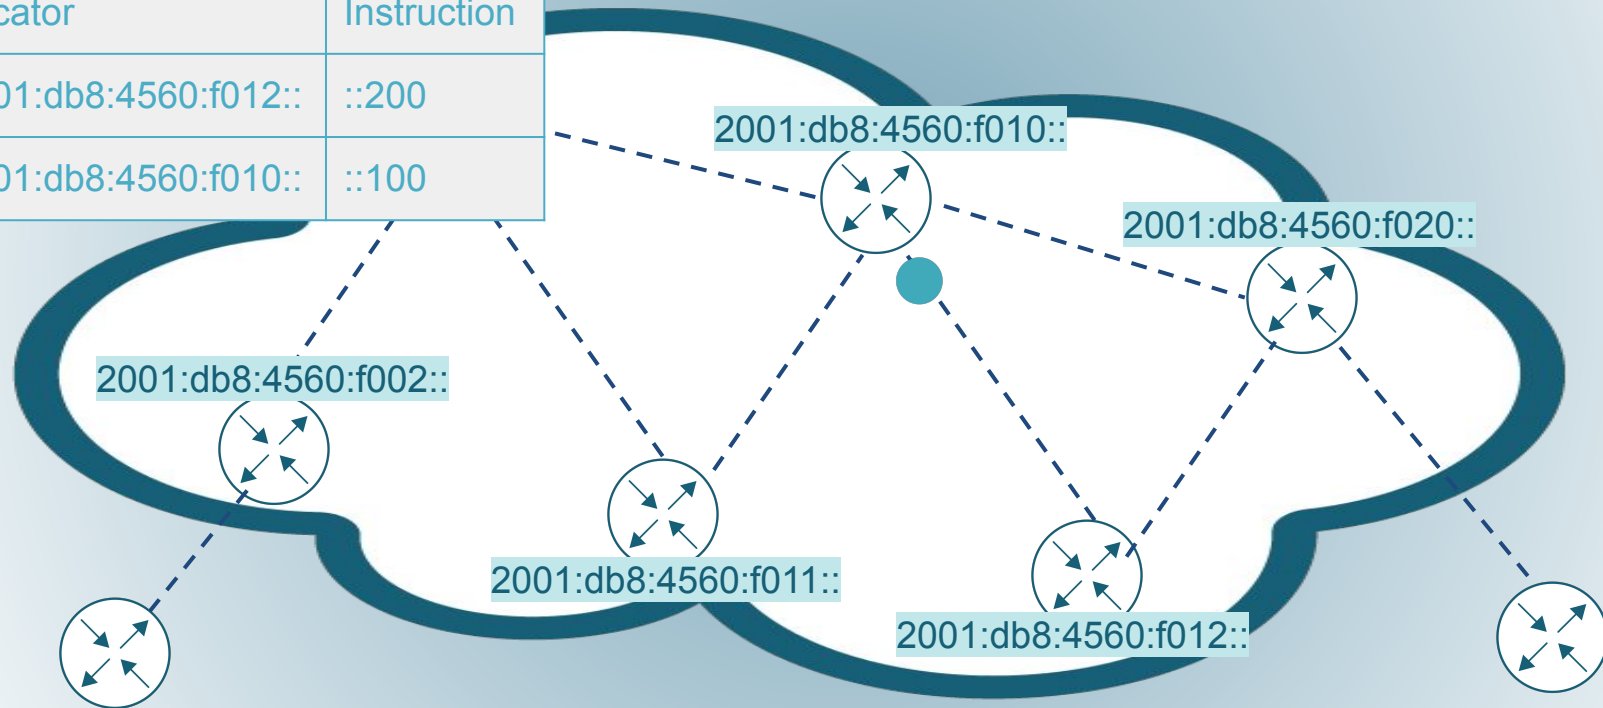
Segment Routing (SRv6)

Locator	Instruction
2001:db8:4560:f010::	::100
2001:db8:4560:f012::	::200
2001:db8:4560:f010::	::100



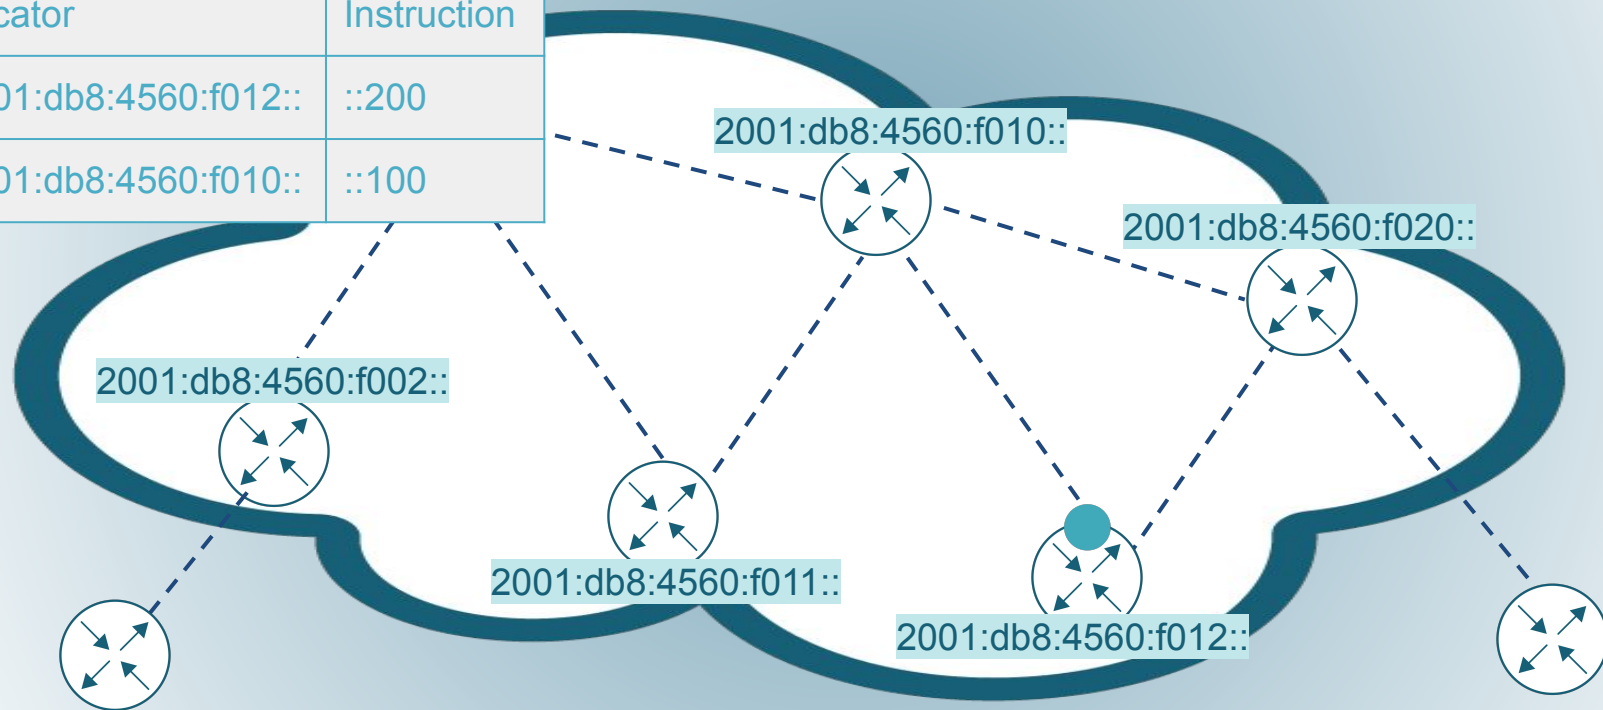
Segment Routing (SRv6)

Locator	Instruction
2001:db8:4560:f012::	::200
2001:db8:4560:f010::	::100



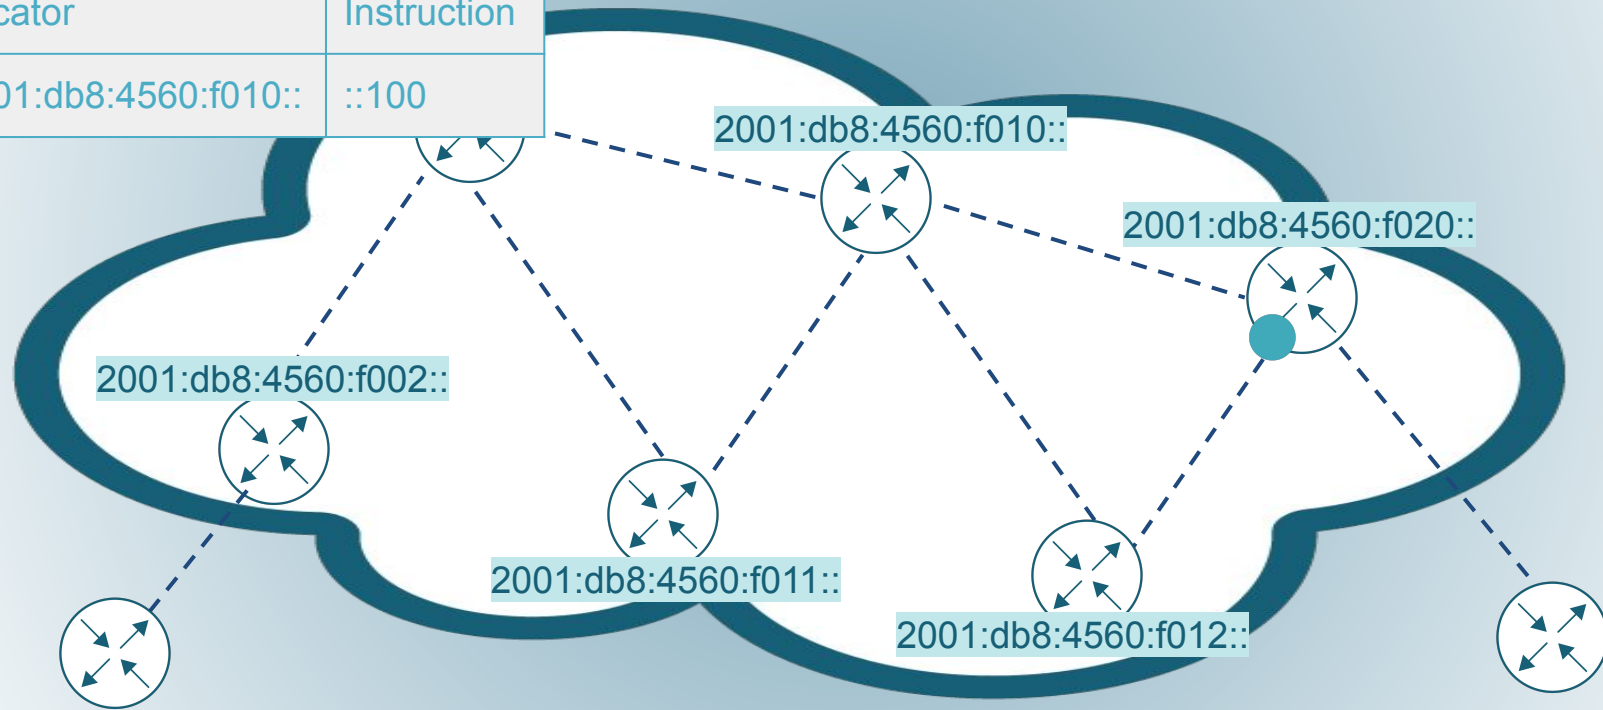
Segment Routing (SRv6)

Locator	Instruction
2001:db8:4560:f012::	::200
2001:db8:4560:f010::	::100



Segment Routing (SRv6)

Locator	Instruction
2001:db8:4560:f010::	::100



SRv6: So What?

- No LDP, RSVP-TE, NSH; underlay and overlay are the same protocol (IP)
- TI-LFA: precalculated backup route for FRR
- Service chaining
 - NFV topology and service are in the same header
 - Chain HW and SW appliances in native IP
- No state tables for NFV or TE
- Incremental deployment
- SDN support implicit

Potential Regulatory Drivers

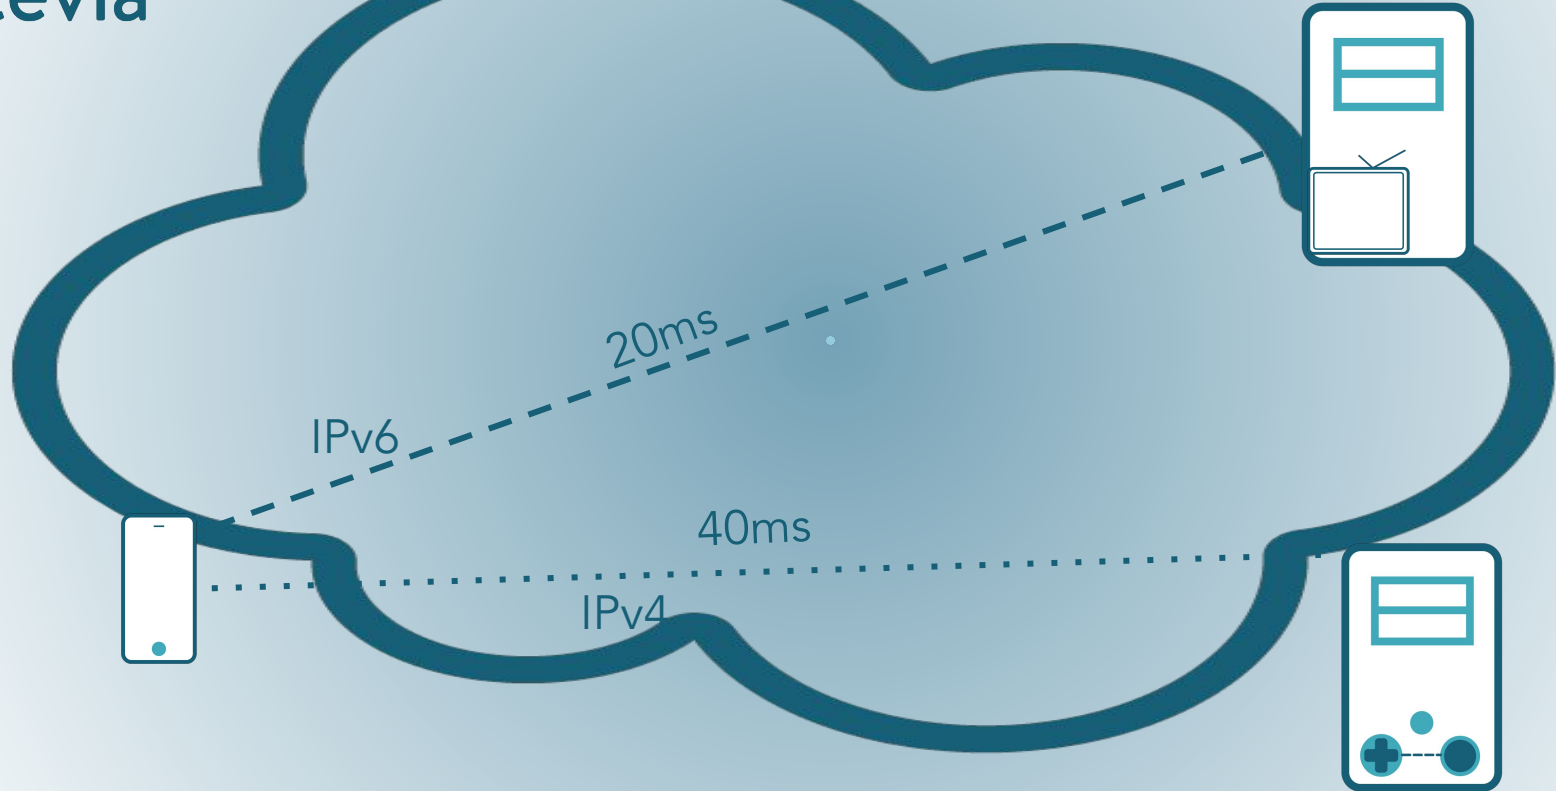
Competition

All I have to do is
outlast the
competition

Why is it so
expensive to
build a new
network?

All I have to do is
outlast the
competition

Neutrality



Perceived Obstacles

Obstacles

- (training, not a priority and why it maybe should be)
- “Lack of customer readiness (55%) and demand (48%) are the main challenges respondents face in relation to IPv6 deployment. A lack of skills and experience within their organisation is also making IPv6 deployment challenging. Reflecting focus group feedback, many organisations also see little economic or operational benefit in implementing IPv6, reducing the urgency to deploy until it is absolutely necessary for their organisation.”

- <https://www.apnic.net/wp-content/uploads/2018/09/2018-APNIC-Member-Survey-Report.pdf>

-

2018 APNIC Member Survey

Our customers are not ready for IPv6	55%
There is no demand for IPv6 from customers	48%
Lack of skills and expertise within our organisation	46%
No clear business / technical advantages or reasons to adopt IPv6	35%
Lack of applications that can run on IPv6	35%
Lack of available training	33%
My organisation's legacy systems do not support IPv6	22%
Our upstream providers do not support IPv6	17%
Cost of IPv6 deployment is too high	16%
The risks of deploying IPv6 are too high	13%

Security

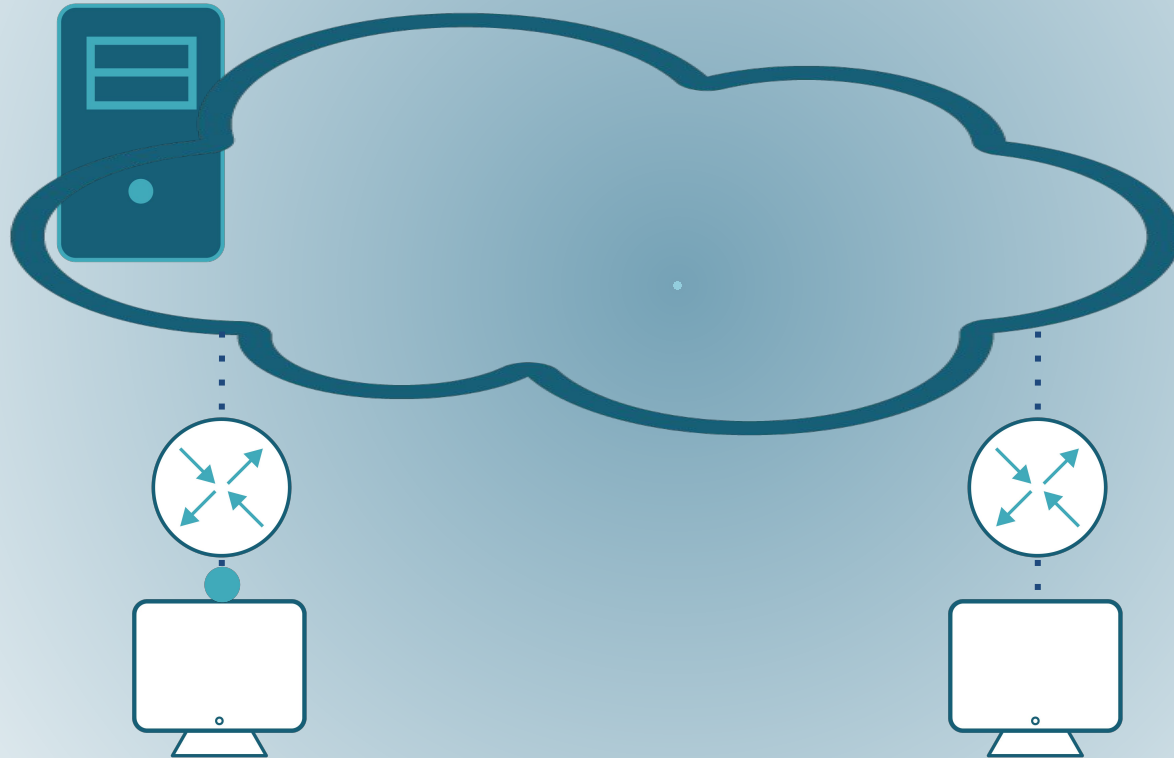
Popular Misconceptions

1. We're safe because we haven't turned on IPv6 yet.
2. NAT keeps us secure.
3. At least with all that address space, host scanning is a thing of the past.
4. IPv6 is more secure because it requires IPSec.

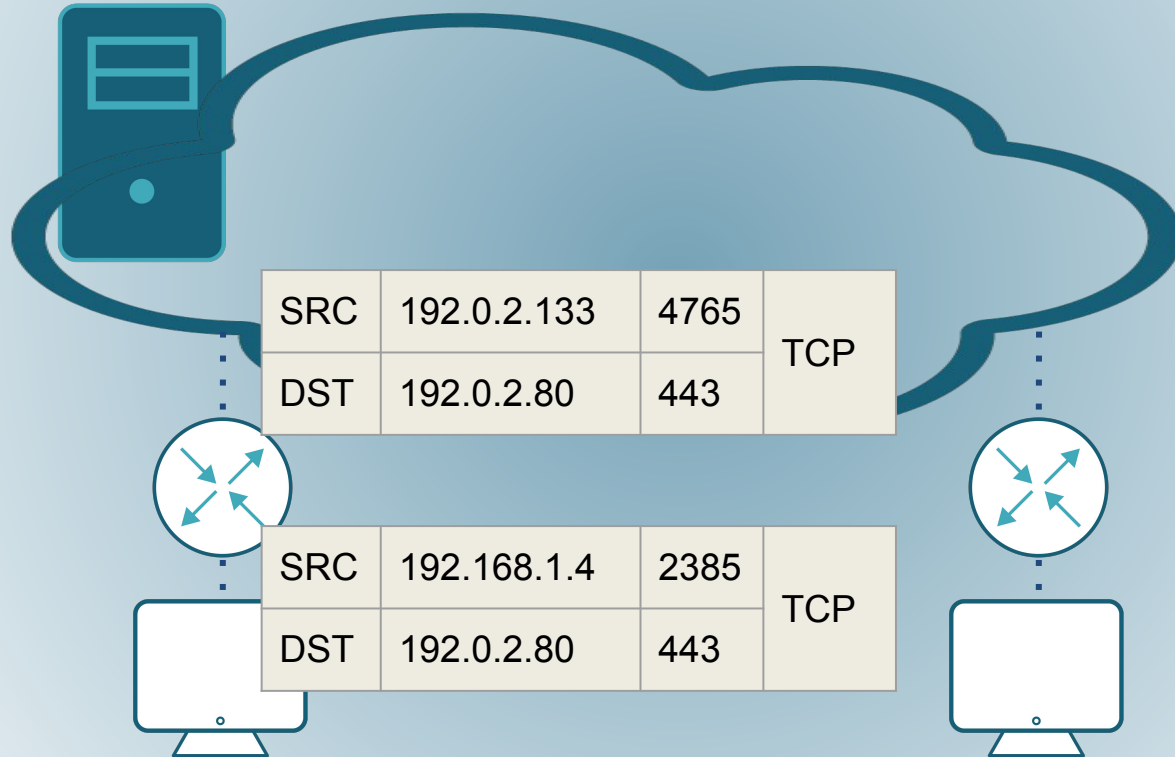
IPv6: On by Default

- Unless you have pushed policies to hosts to disable IPv6, LLA is already turned on
- Some firewalls have IPv6 open by default
- Some IDS/IPS ignore unrecognized traffic
- Many IPv6 transition technologies are tunnels

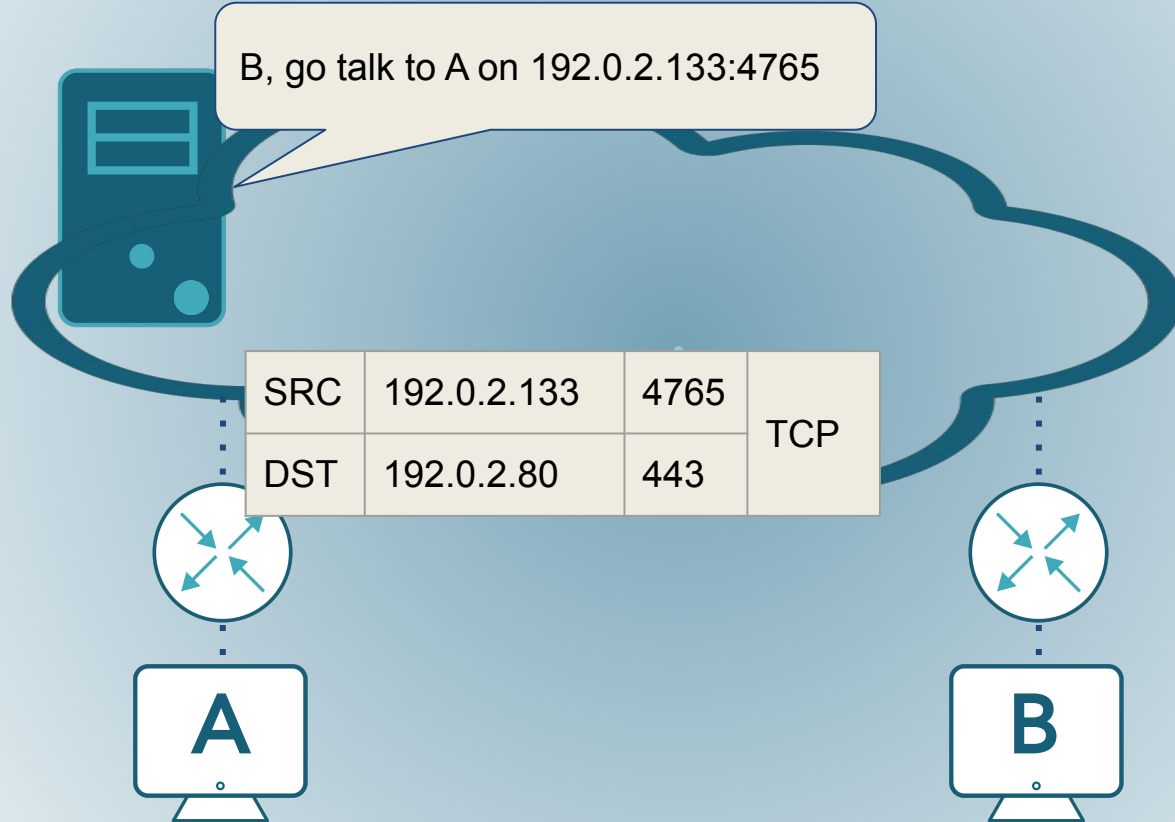
NAT is not a Firewall



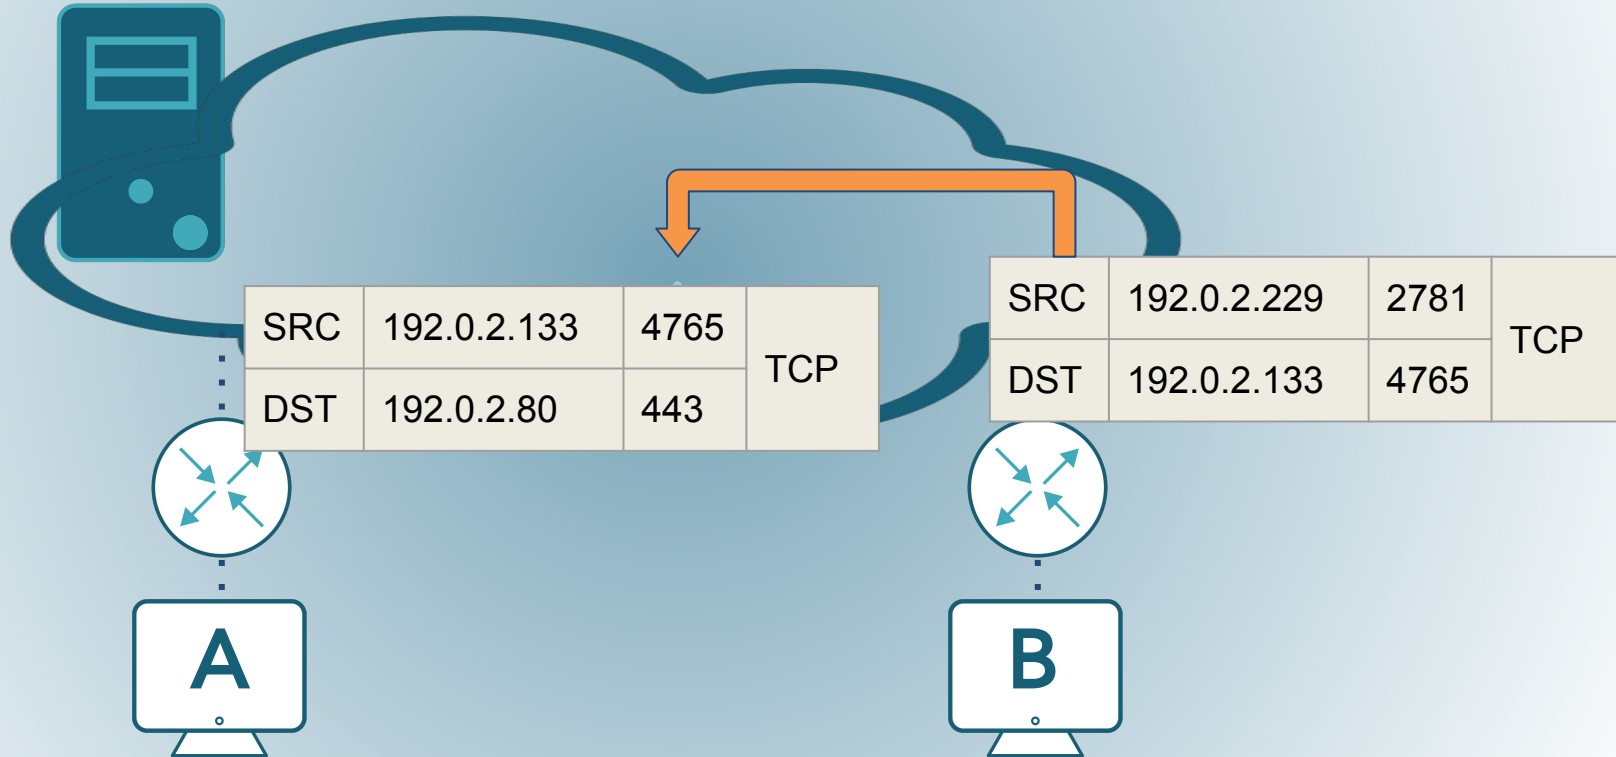
Basic NAT Translation



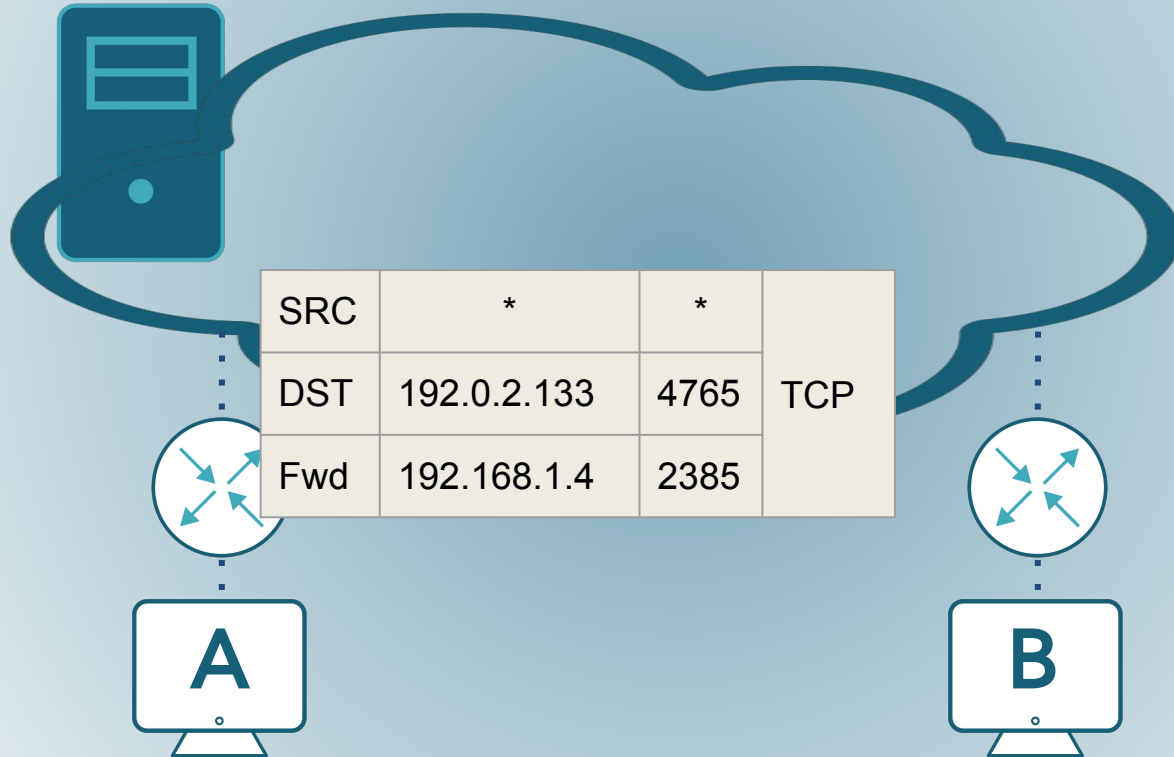
What about p2p or gaming?



If NAT was FW, packet drops



Full cone NAT forwards *



Host Scanning

- $2^{64} = 18,446,744,073,709,551,616$ addresses
- But within `2001:db8:f001:1::/64` likely host addresses include
 - `::1`
 - `::2`
 - `::80`
 - `::1:1`
 - `::beef`
 - `::<192.0.2.x>`

Host Scanning

- 2001:db8:f001:1::/64 where host bits are EUI-64
 - ::<OUI>ff:feXX:XXXX
 - Pick OUIs from popular NICs and scan 16M addresses
- Lookup or xfer DNS and rDNS
 - Q 1.0.8.0.0.8.b.d.0.1.0.0.2.ip6.arpa.
 - NODATA means the zone exists, so scan for hosts
 - NXDOMAIN means no zone, probably no hosts
- Scan BitTorrent sites or other servers for address logs
-

Host Scanning Mitigations

- FW/IPS blocking ICMPv6 that looks like scanning
- FW or host configured to drop ICMPv6 Echo Request
 - But not ICMPv6 PTB!
 - Policing is possible to prevent DoS of large packet floods,
 - But too-big packets can only arrive on routers with links of different MTUs
- Ignore what I said earlier about mnemonic addresses
- Privacy extensions: randomly change address

IPSec will save us!

Rfc2401 "Security Architecture for the Internet Protocol"
says

```
This section defines Security Association management requirements for  
all IPv6 implementations and for those IPv4 implementations that  
implement AH, ESP, or both.
```

So it's mandatory!

LOCAL RISKS



Vulnerability

- Unauthenticated ND, RA, etc. (same as ARP)
 - Hello, I'm 2001:db8::1
 - No, I'm 2001:db8::1
 - Hello, I'm a router for 2001:db8::/32
- Cache table exhaustion

SLAAC vs DHCPv6

- Some admins like DHCP because it logs who has what address
 - Except it doesn't prevent manual configuration
- Mitigations for rogue attachments
 - Log Neighbor Discovery tables
 - Syslog, SNMP, Netconf
 - 802.1x

Smurf

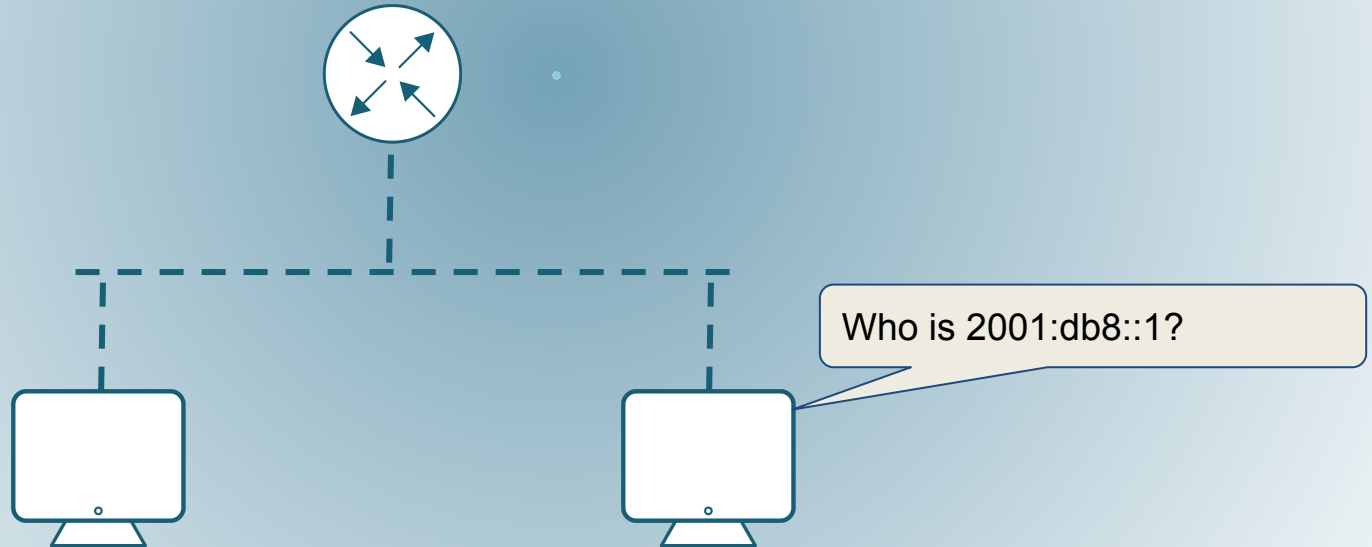
Send packets with spoofed source address (the victim) to a multicast address, for many responses to DOS the victim

Address	Description	Scope
FF01::1	All Nodes Address	Node-Local
FF01::2	All Routers Address	Node-Local
FF02:0:0:0:0:0:1	All Nodes Address	Link-Local
FF02:0:0:0:0:0:2	All Routers Address	Link-Local
FF02:0:0:0:0:0:5	OSPFv2	Link-Local
FF02:0:0:0:0:0:6	OSPFv2 Designated Routers	Link-Local
FF02:0:0:0:0:0:C	SSDP	Link-Local
FF02:0:0:0:0:0:12	VRRP	Link-Local
FF02:0:0:0:0:0:FB	mDNSv6	Link-Local
FF02:0:0:0:0:0:1:2	All_DHCP_Relay_Agents_and_Servers	Link-Local

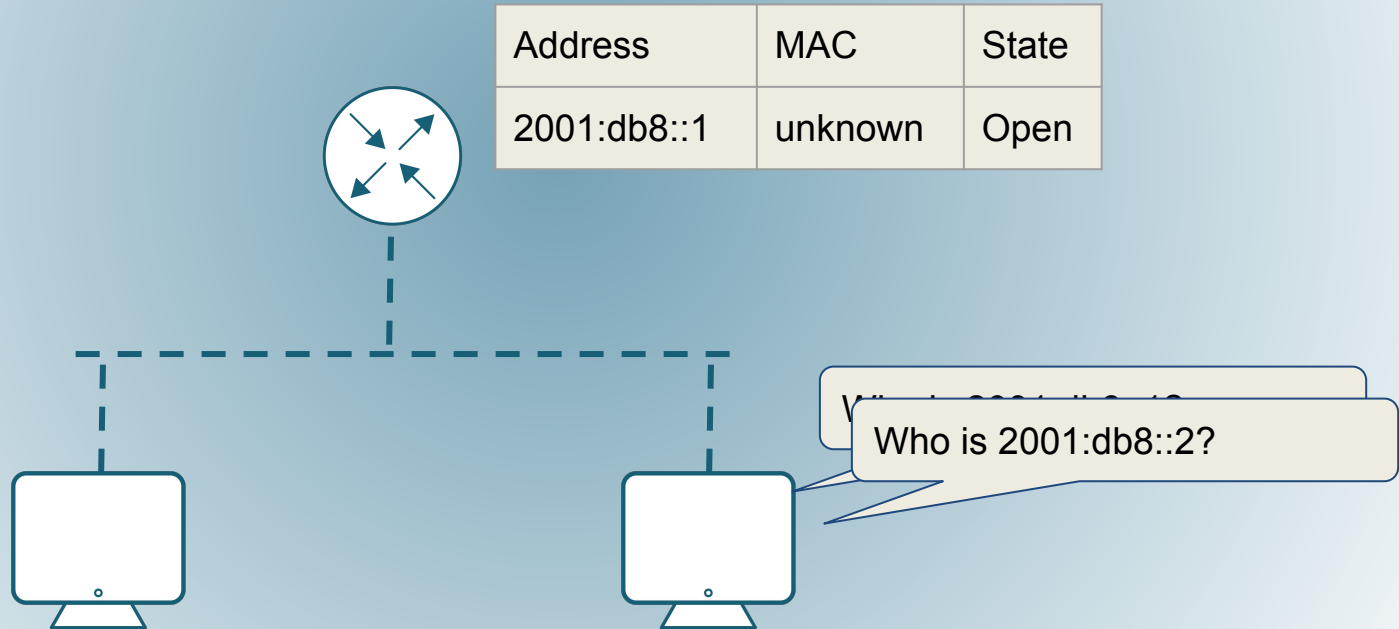
And many more!

<https://www.iana.org/assignments/ipv6-multicast-addresses/ipv6-multicast-addresses.xhtml>

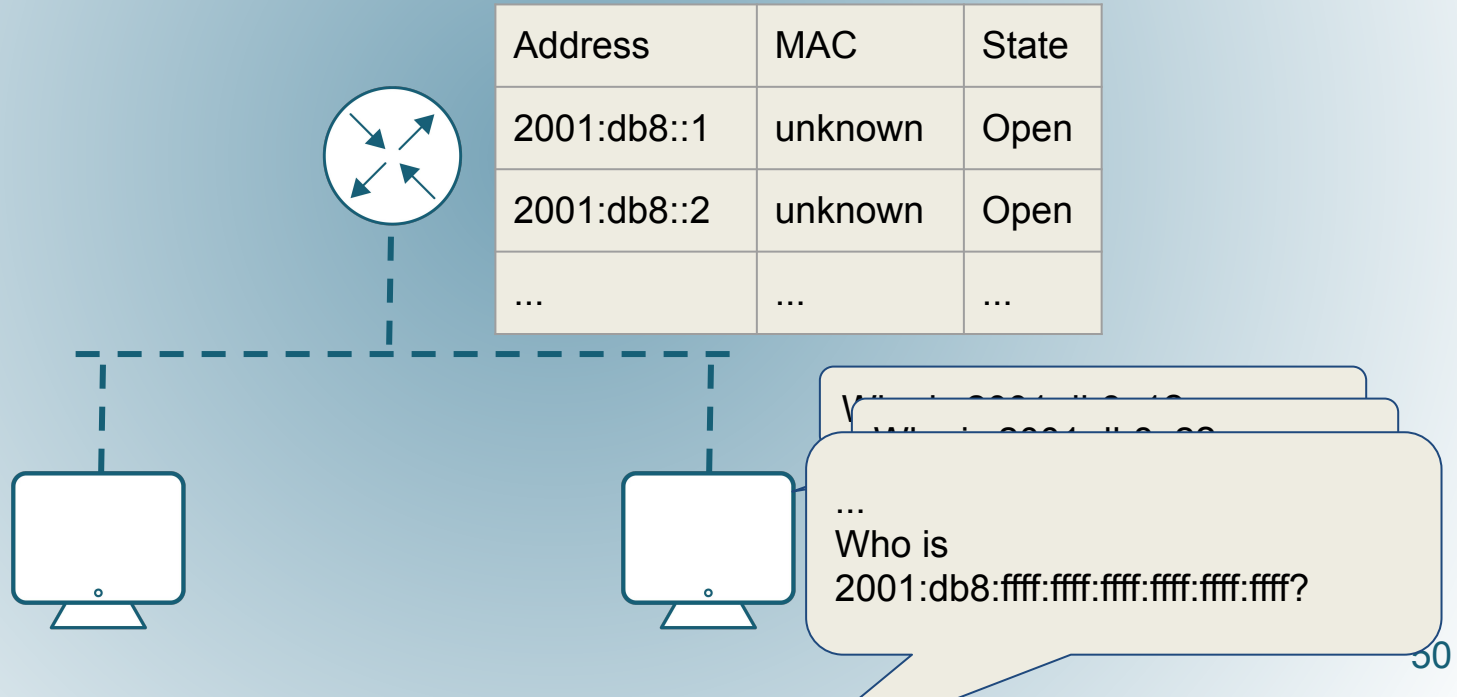
Neighbor Table Exhaustion



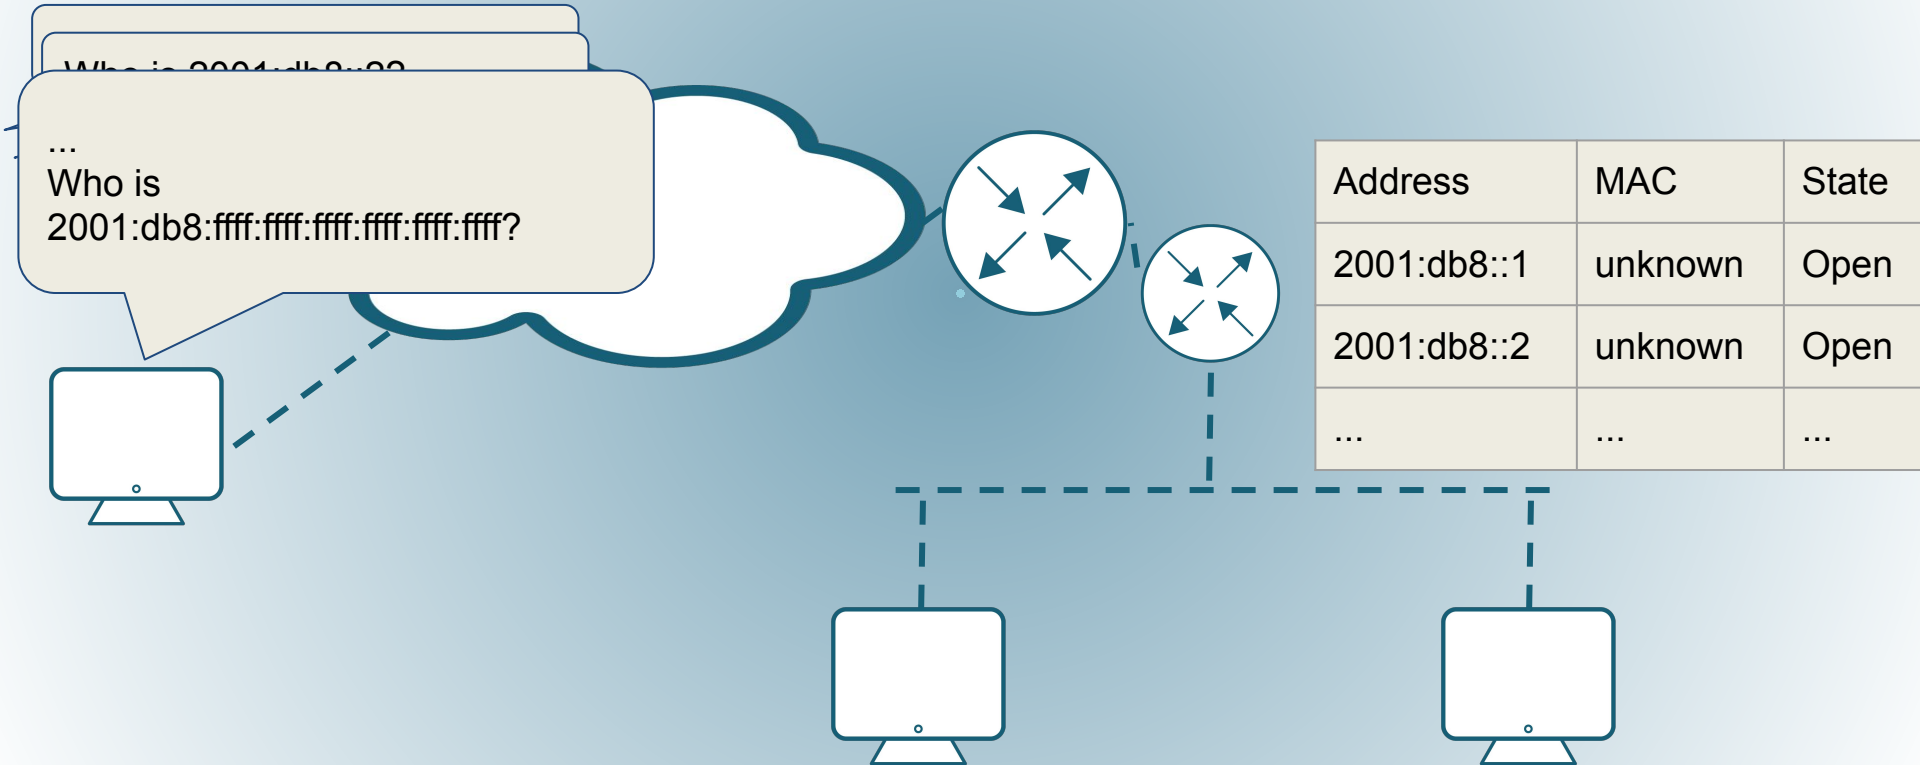
Neighbor Table Exhaustion



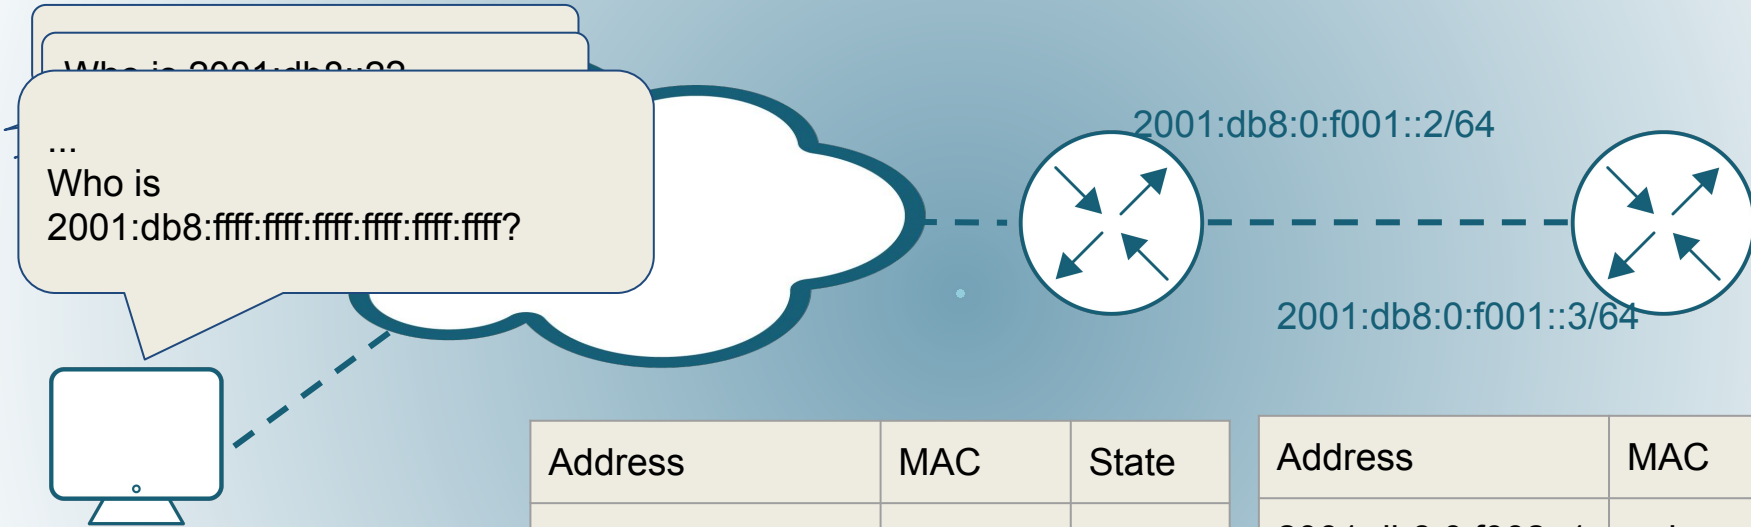
Neighbor Table Exhaustion



Neighbor Table Exhaustion



Ping Pong Attack



Address	MAC	State
2001:db8:0:f002::1	unknown	Open
2001:db8:0:f002::4	unknown	Open
...

Address	MAC	State
2001:db8:0:f002::1	unknown	Open
2001:db8:0:f002::4	unknown	Open
...

NDT Mitigations

- /127 netmask
- ACL on unused space
- NDP Queue rate limit
 - If device has different queues for confirming existing entries and resolving new queries, tighten new query queue
- Rate limit ICMPv6
- and several mechanisms to log bad NDP. . .

- Secure path to CA
 - Send request for CA
 - Each node on the path sends its cert
 - CA confirms each cert
- Use key pair to generate CGA
 - CryptoGraphically Assigned host bits
- Send RS; Router replies with signed RA
- Uses SHA-1 and PKIX; not highly secure
 - Because longer keys would exceed MTU, requiring frag

RA-Guard

- L2 switch can prevent malicious/spurious RAs
- Multiple possible policies
 - Block RAs from specific MAC or port
 - Allow RAs only from specific MAC or port
 - Allow RAs that comply with (e.g., SeND) policy
 - Or use prefix list, prefix range, router priority
- Switch can become RA proxy
- Off -> Learning -> Blocking -> Forwarding

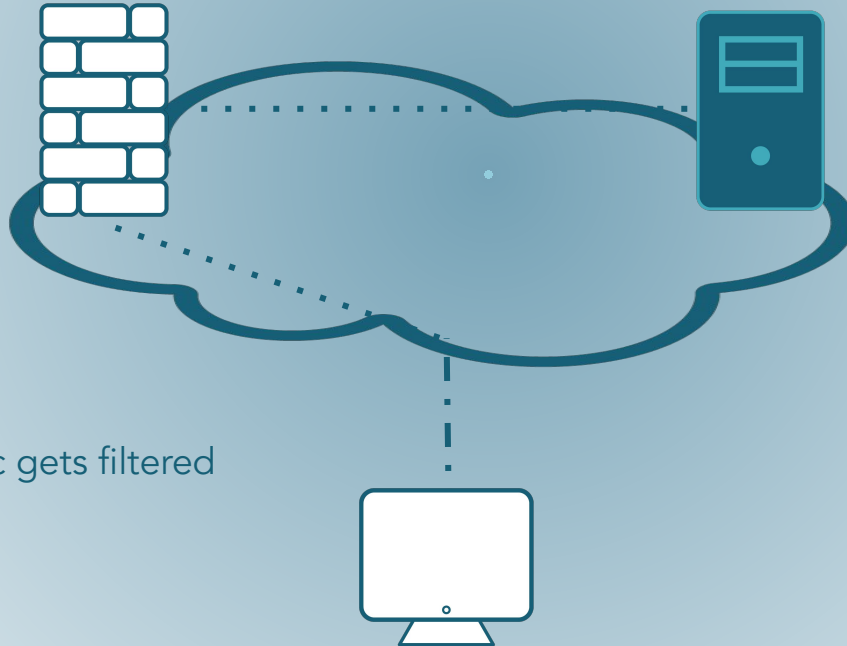
SAVI

- Source Address Verification Improvements against spoofing
- FCFS SAVI: first user of address (within prefix list or RA) is authorized user
- SeND SAVI: drop packets where SRC not certified
- SAVI with DHCP: snoop DHCP, drop packets from IP addresses not assigned by DHCP
- SAVI-MIX: if two SAVIs conflict, resolve in order

Cisco, in their IPv6-only enterprise network

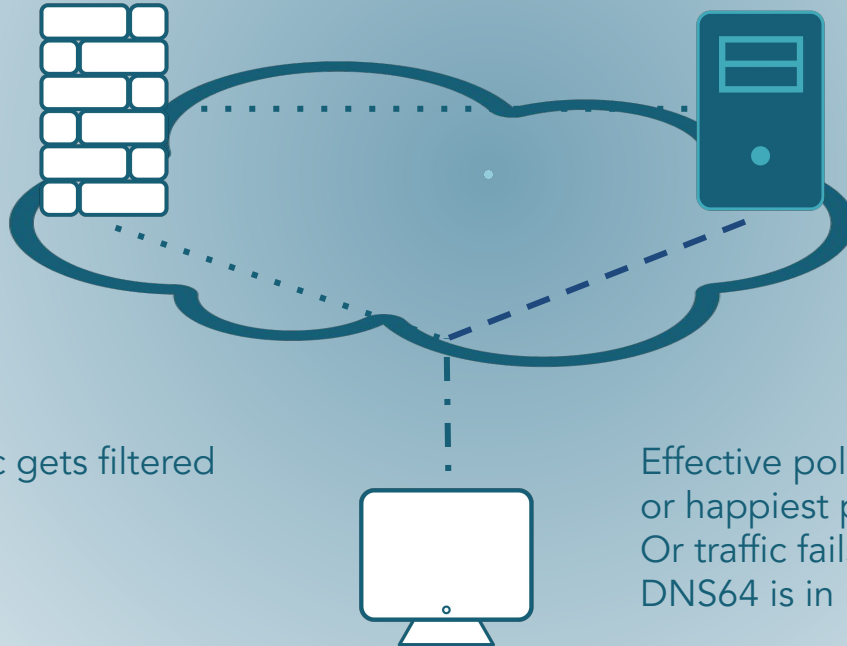
- First Hop Security
 - IPv6 Snooping (Address Gleaning, Device Tracking)
 - ND Inspection
 - DHCPv6 Guard
 - RA Guard
 - Source Guard
- Data center: Cloud security

VPN



Intended policy: traffic gets filtered through VPN

VPN



Intended policy: traffic gets filtered through VPN

Effective policy: traffic takes shortest or happiest path
Or traffic fails to DS server, or if DNS64 is in use

Fragmentation

- Remember that only sender can fragment
- SeND RA might be too big and require frag
 - Local sender could send fragments that collide with SeND
- RA with many PIOs might require frag
 - Send multiple RAs instead
- Good place to troubleshoot if RAs are failing silently

FIREWALL SPECIFICS



Extension Headers

- Extension Headers
 - HBH
 - DO
 - Routing
 - Fragment
 - AH, ESP
 - Others. . . see IANA registry
- L4 or higher inspection?
 - Parse all headers to find pointer to the Upper-Layer Header

ICMPv6

- Link local multicast and address discovery
- ICMPv6 message types
 - Destination Unreachable
 - PTB
 - Time Exceeded
 - Parameter Problem
 - Echo Request
 - Echo Reply

Spam

- 22/50 top sites have IPv6 MX records
 - 20 of them use Google for mail.
 - LinkedIn, WikiMedia.
- IP reputation tools are terrible at IPv6
 - Block /64? /60? /56? /48?

IPv6-Specific Security Tools

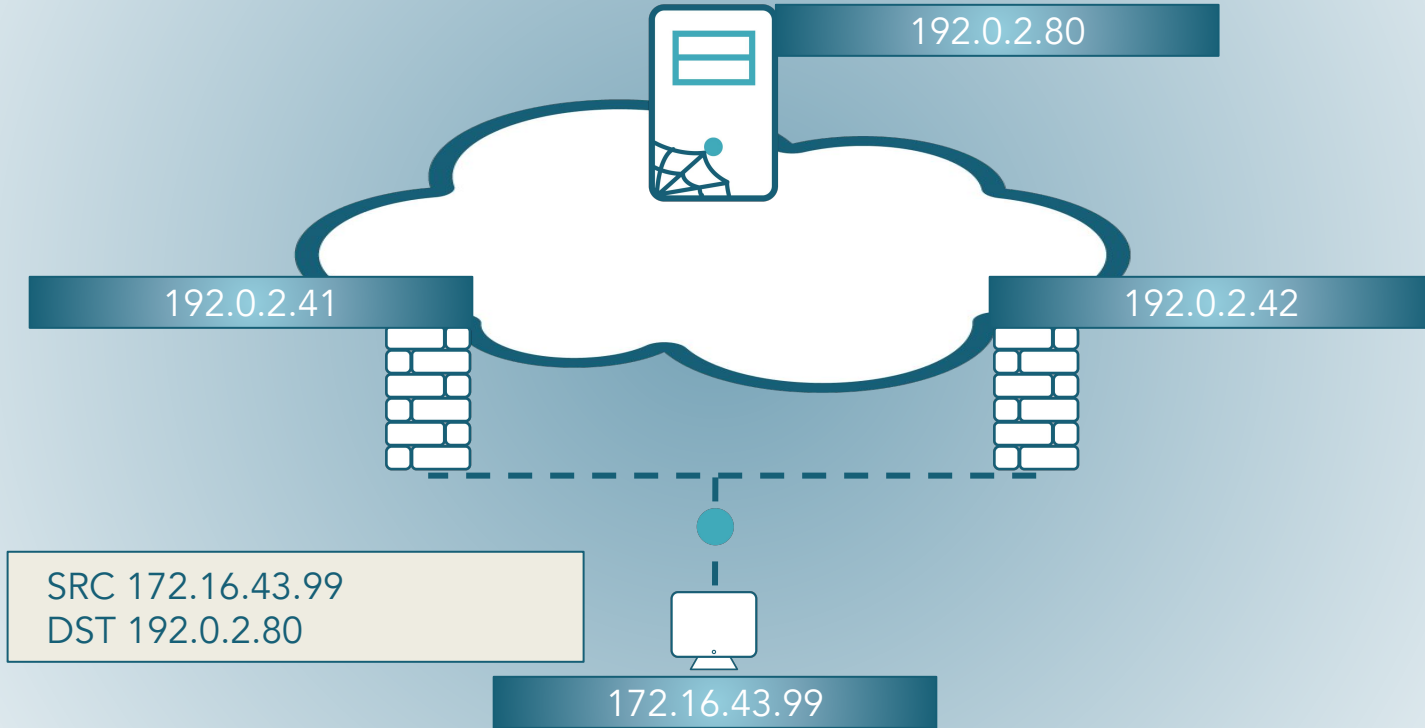
- THC
- IPv6-Toolkit
- FT6 Firewall Tester
- Many existing tools

Running a dual-stack network doubles the attack exposure as a malevolent person has now two attack vectors: IPv4 and IPv6.

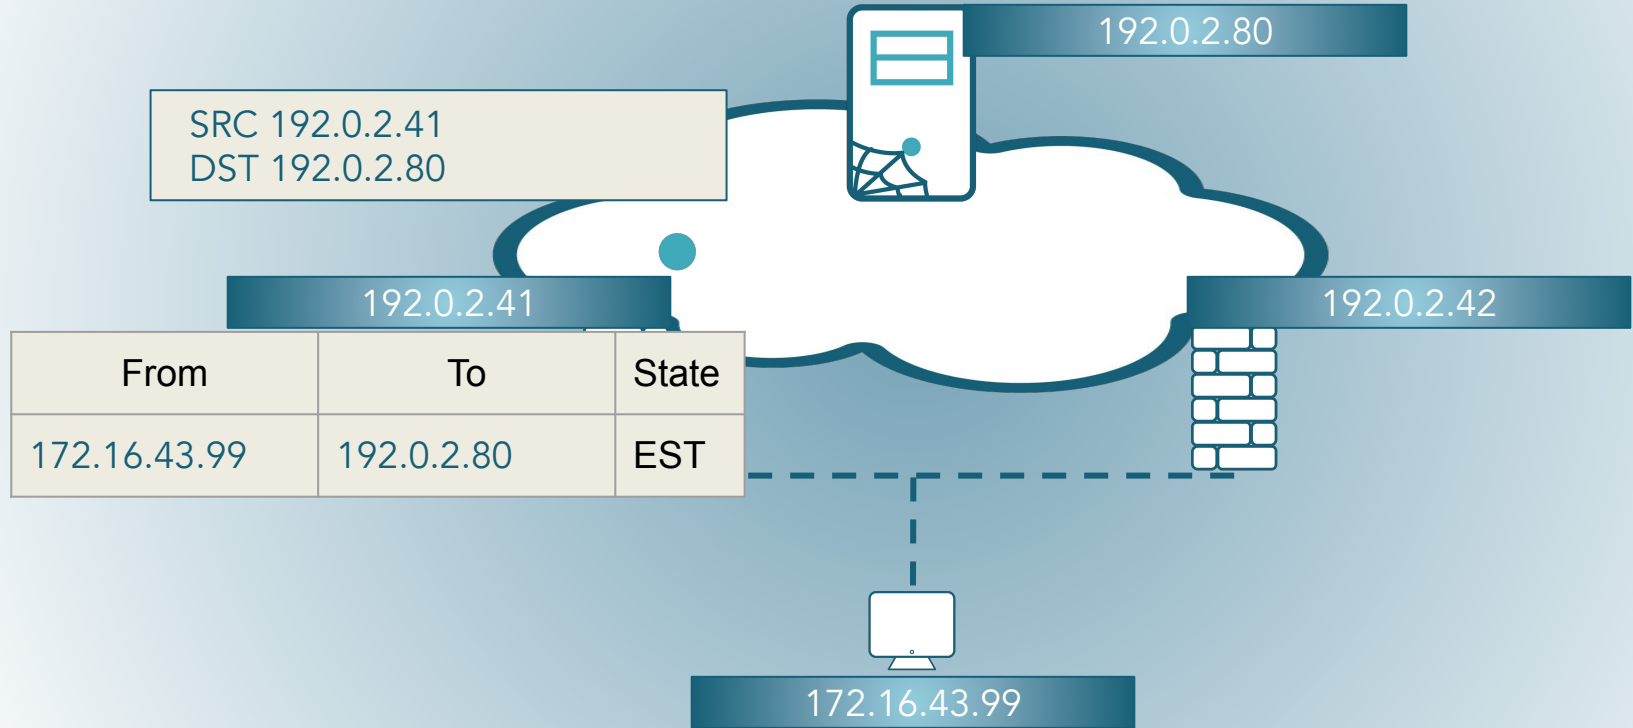
--RFC7381 "Enterprise IPv6 Deployment Guidelines"

The Multihoming Problem

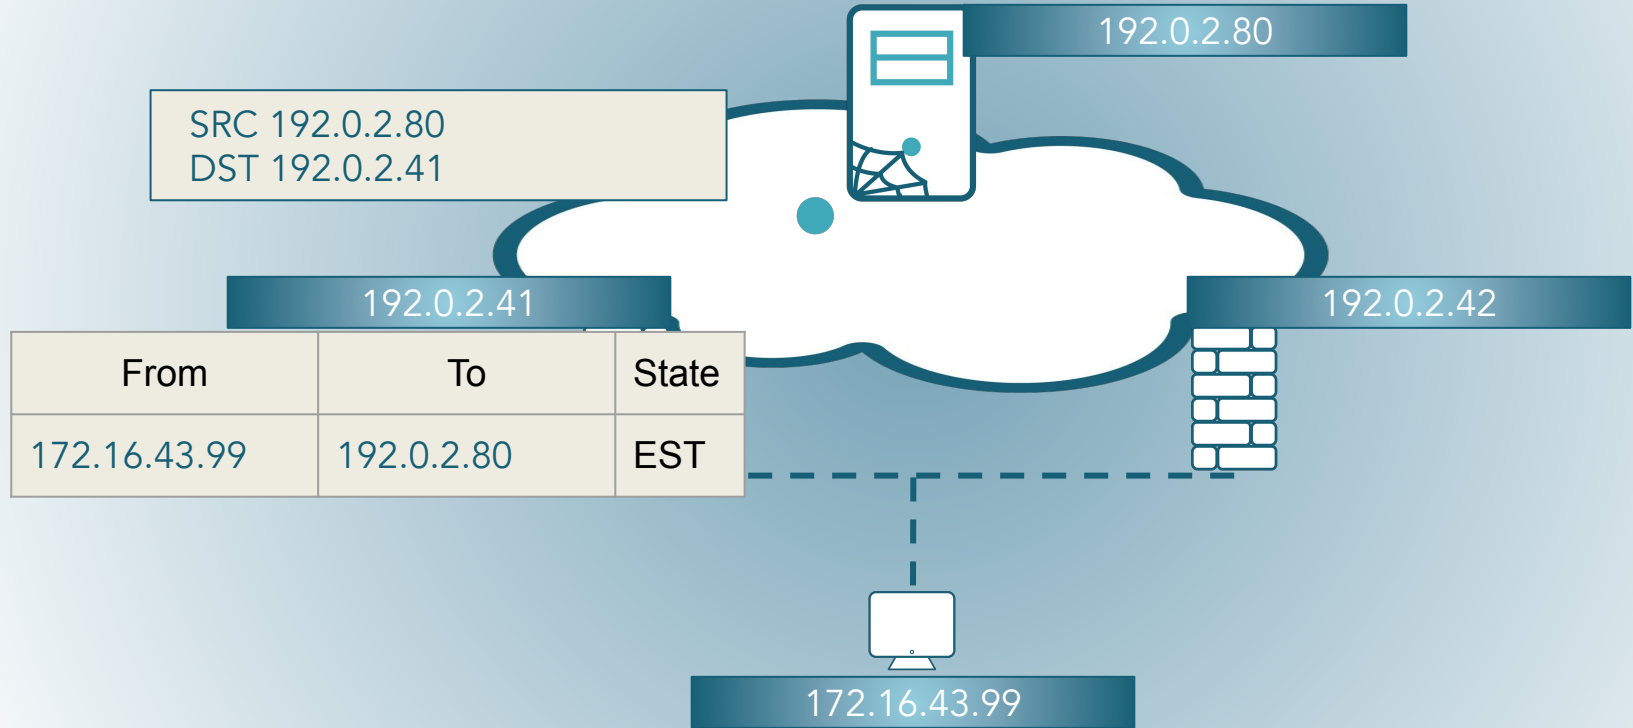
Multihoming Status Quo



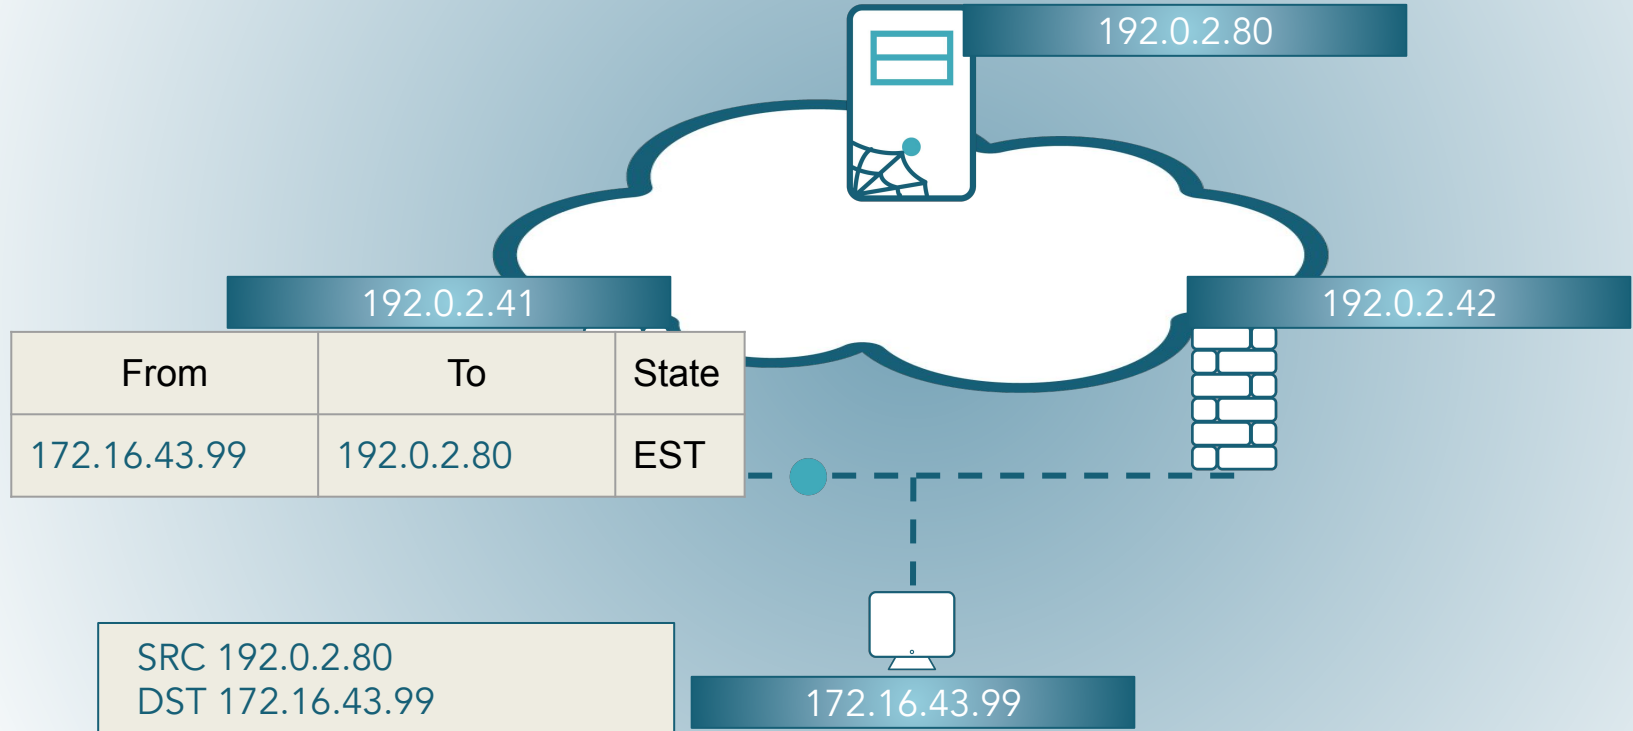
Multihoming Status Quo



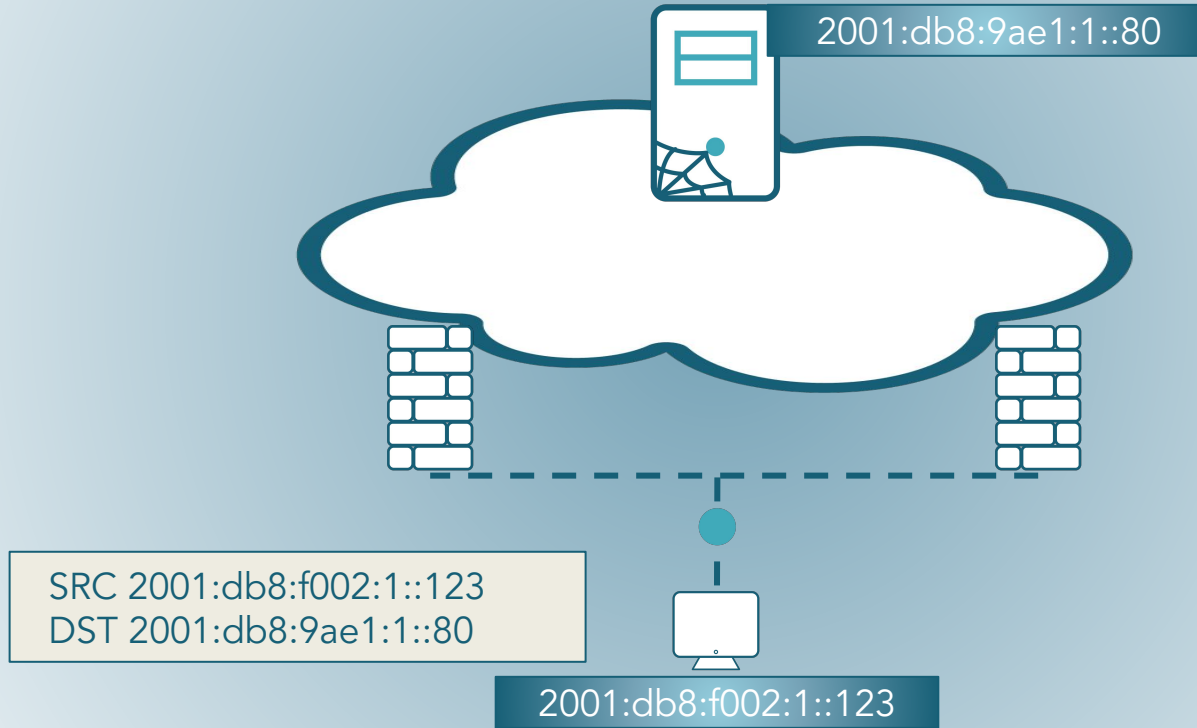
Multihoming Status Quo



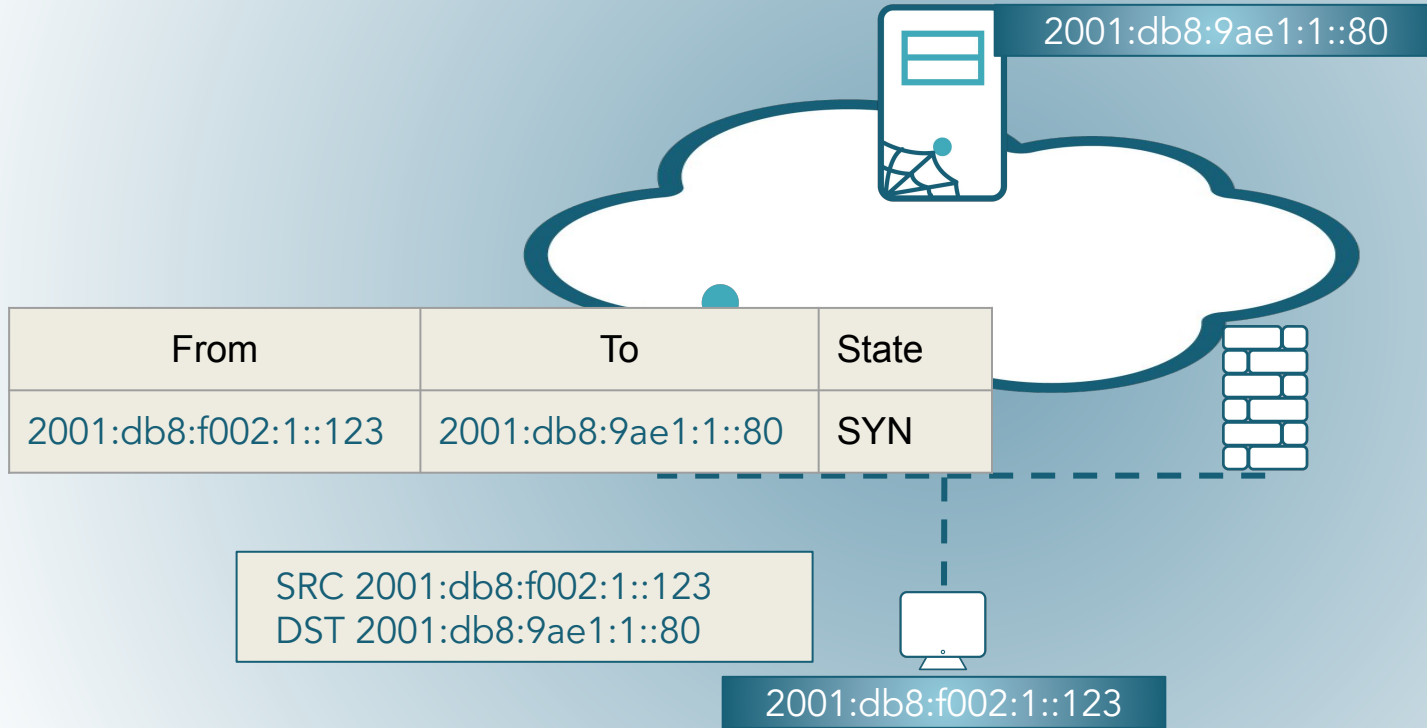
Multihoming Status Quo



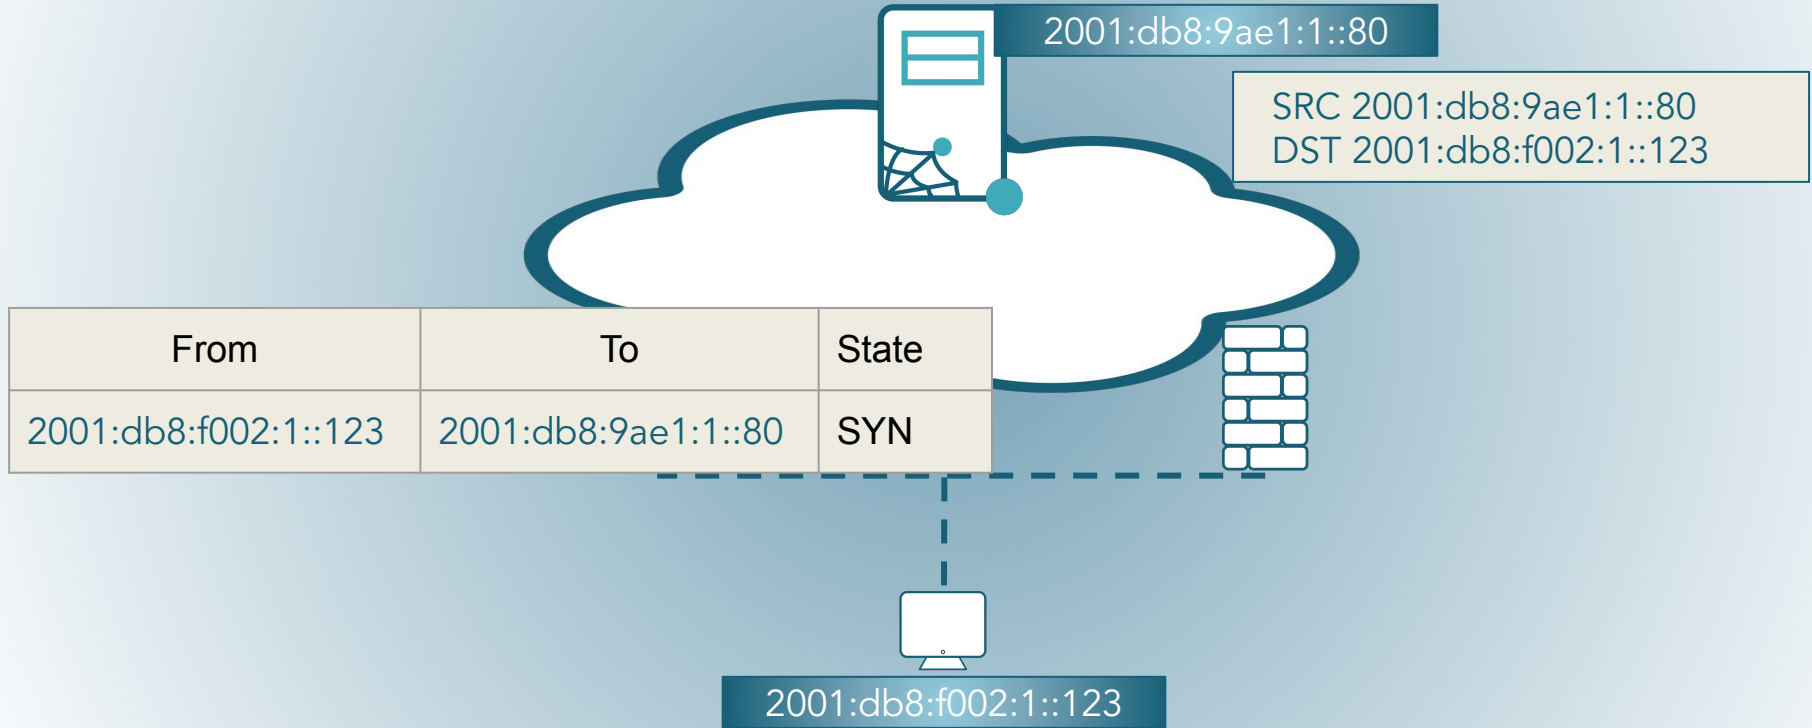
The Multihoming Problem



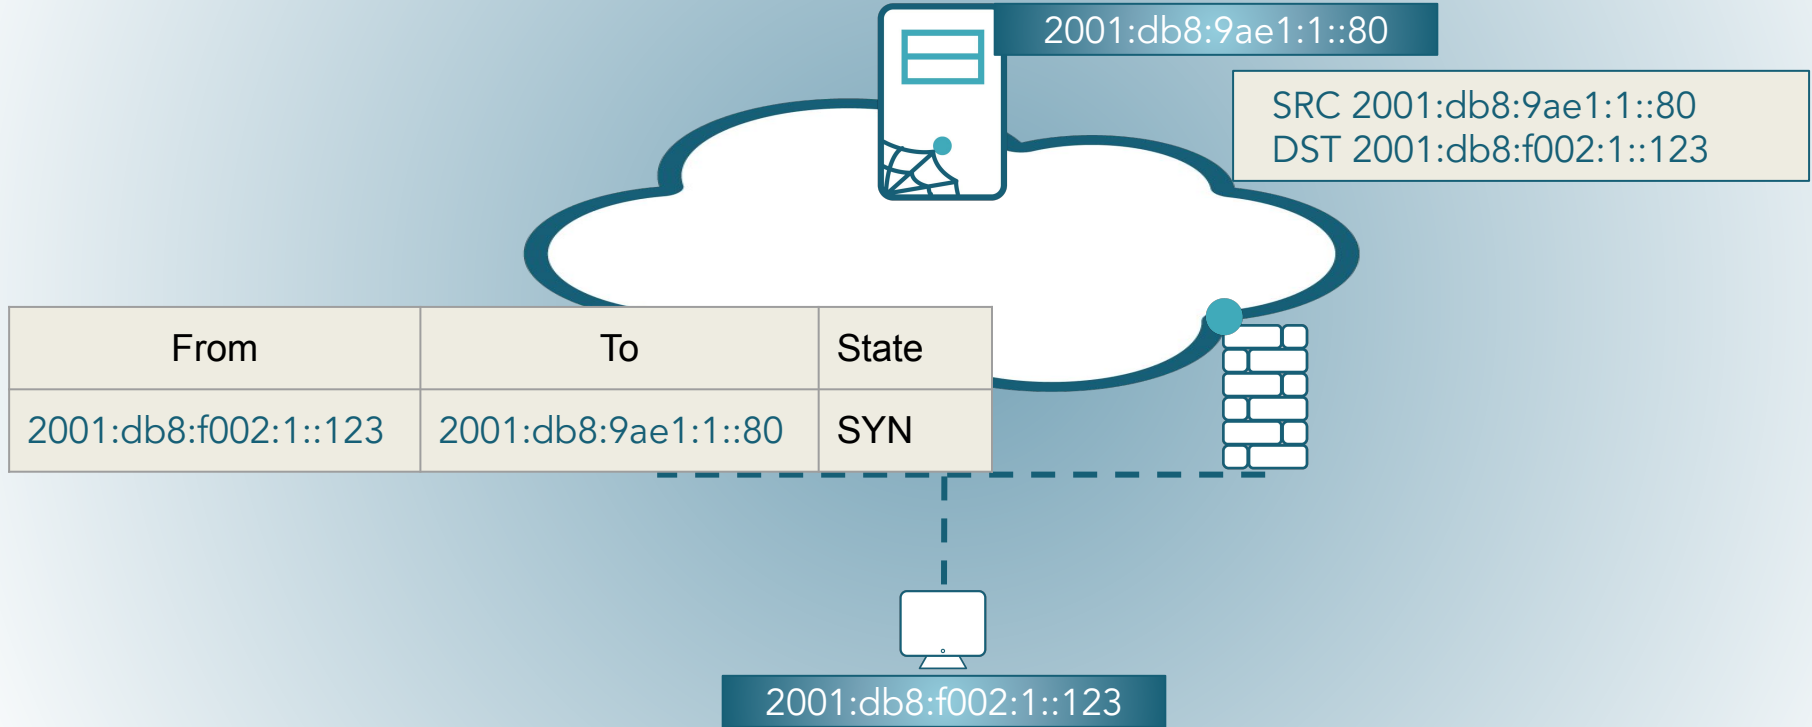
The Multihoming Problem



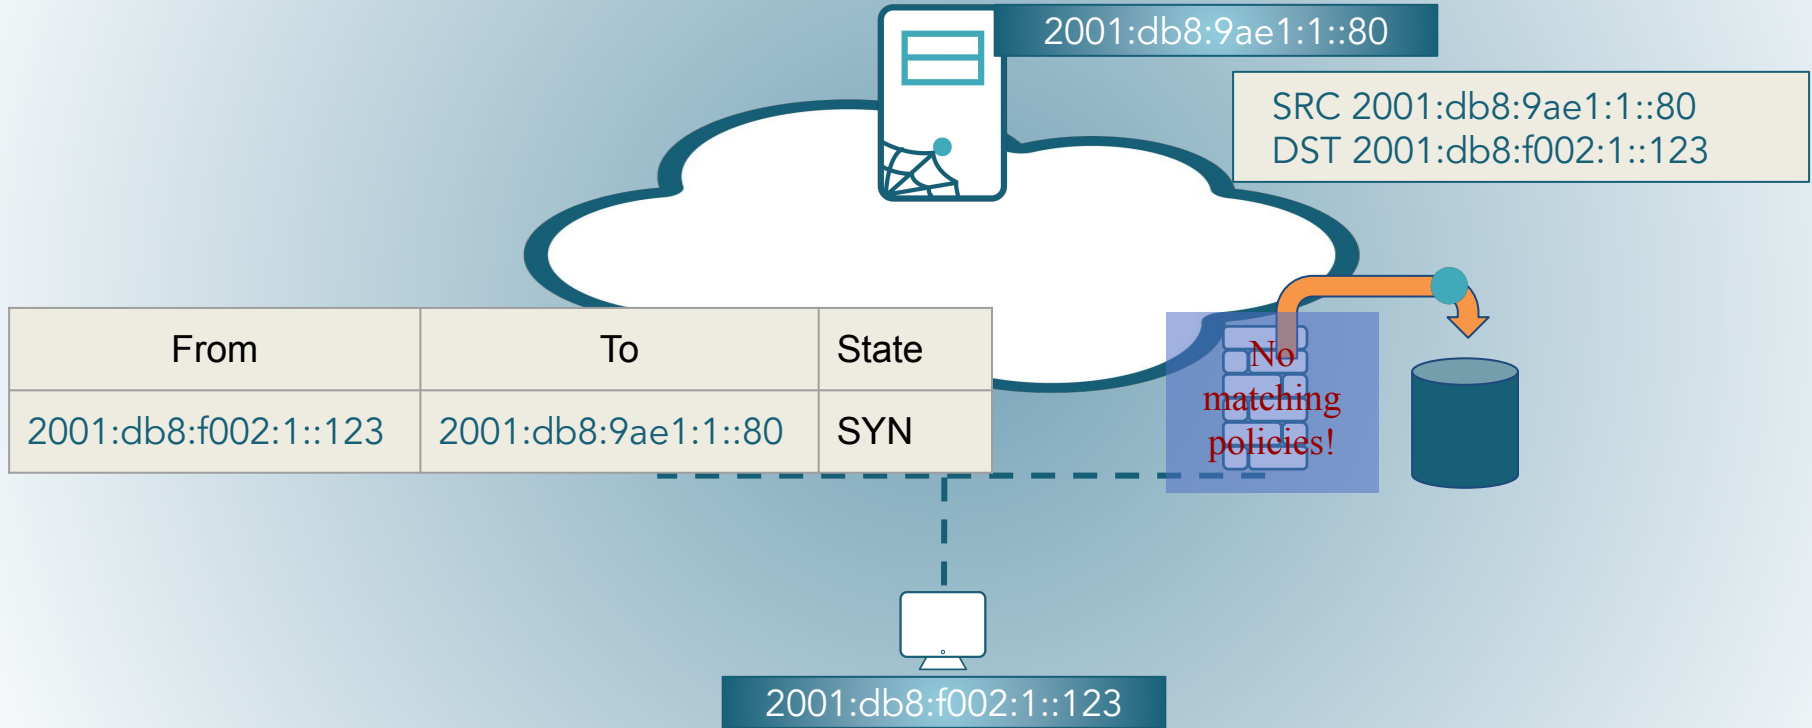
The Multihoming Problem



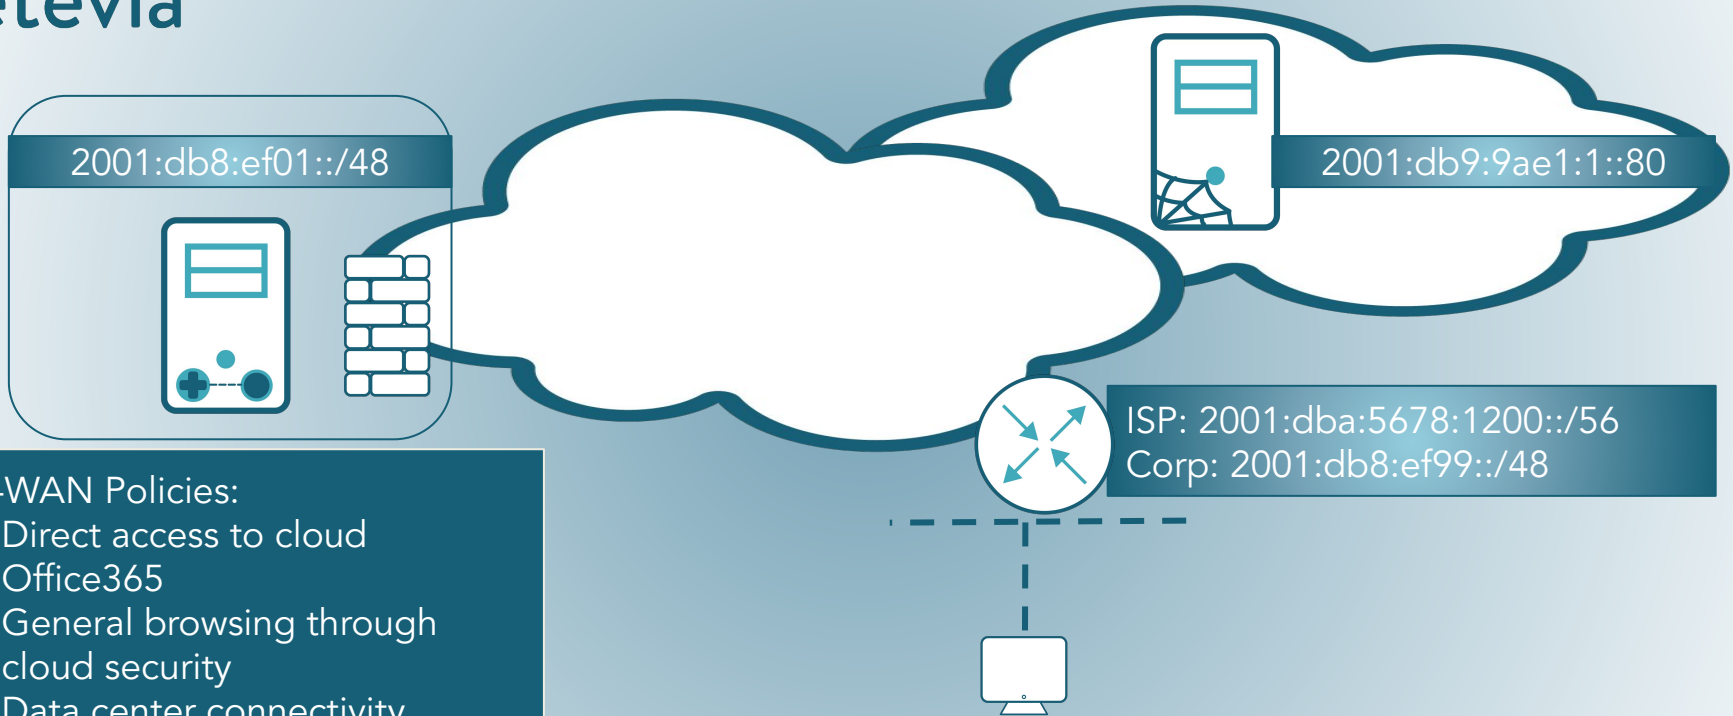
The Multihoming Problem



The Multihoming Problem



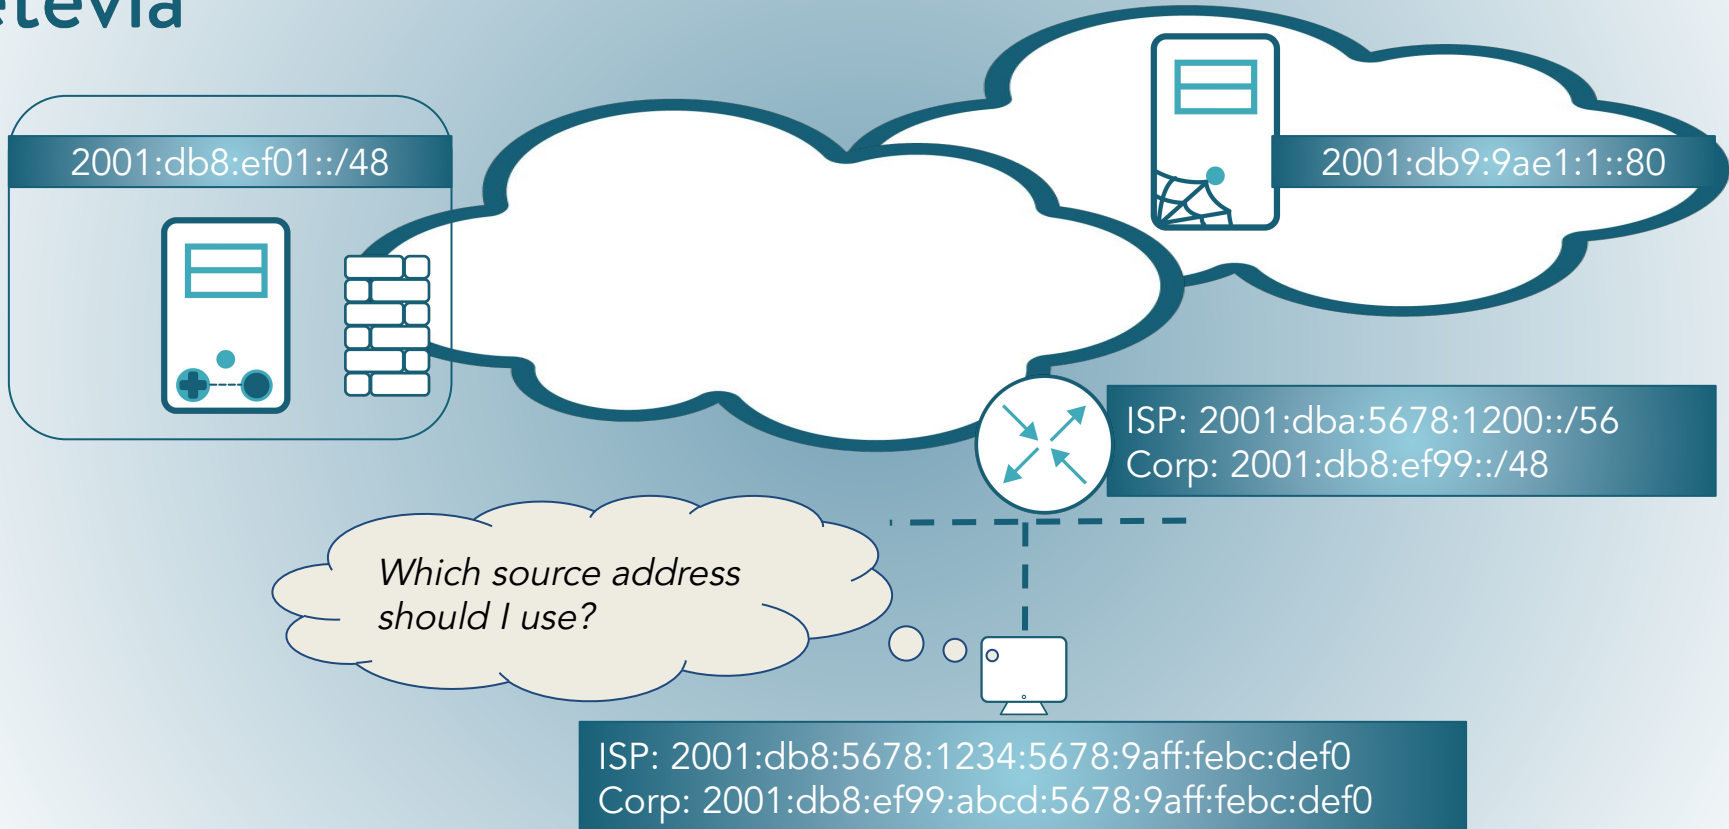
Two Networks



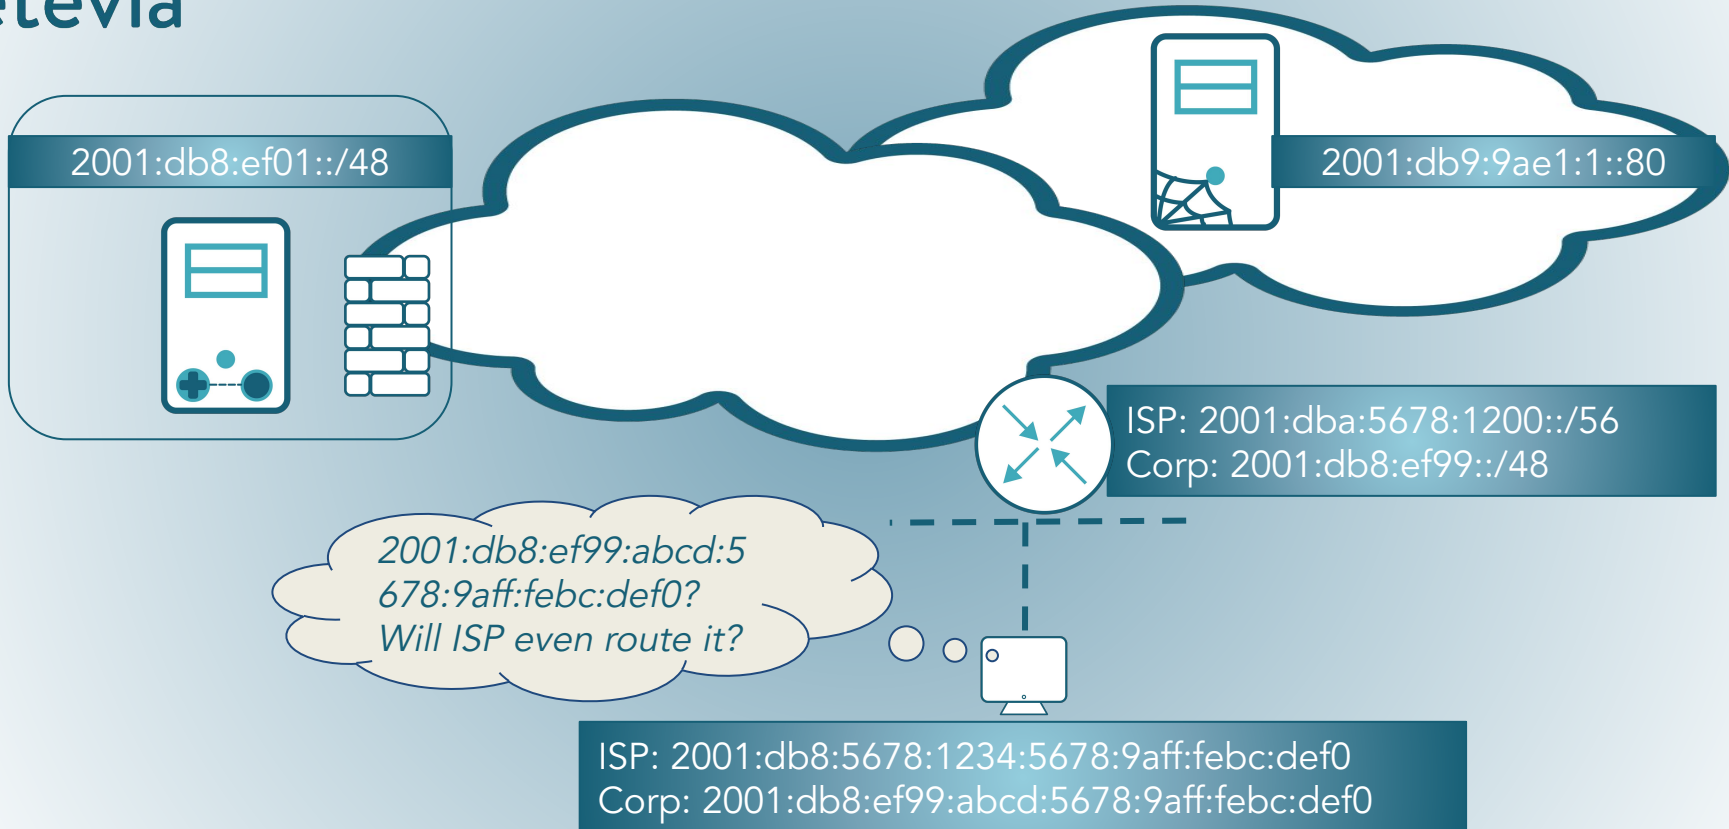
SD-WAN Policies:

1. Direct access to cloud Office365
2. General browsing through cloud security
3. Data center connectivity through MPLS VPN
 - a. Backup option through encrypted Internet VPN

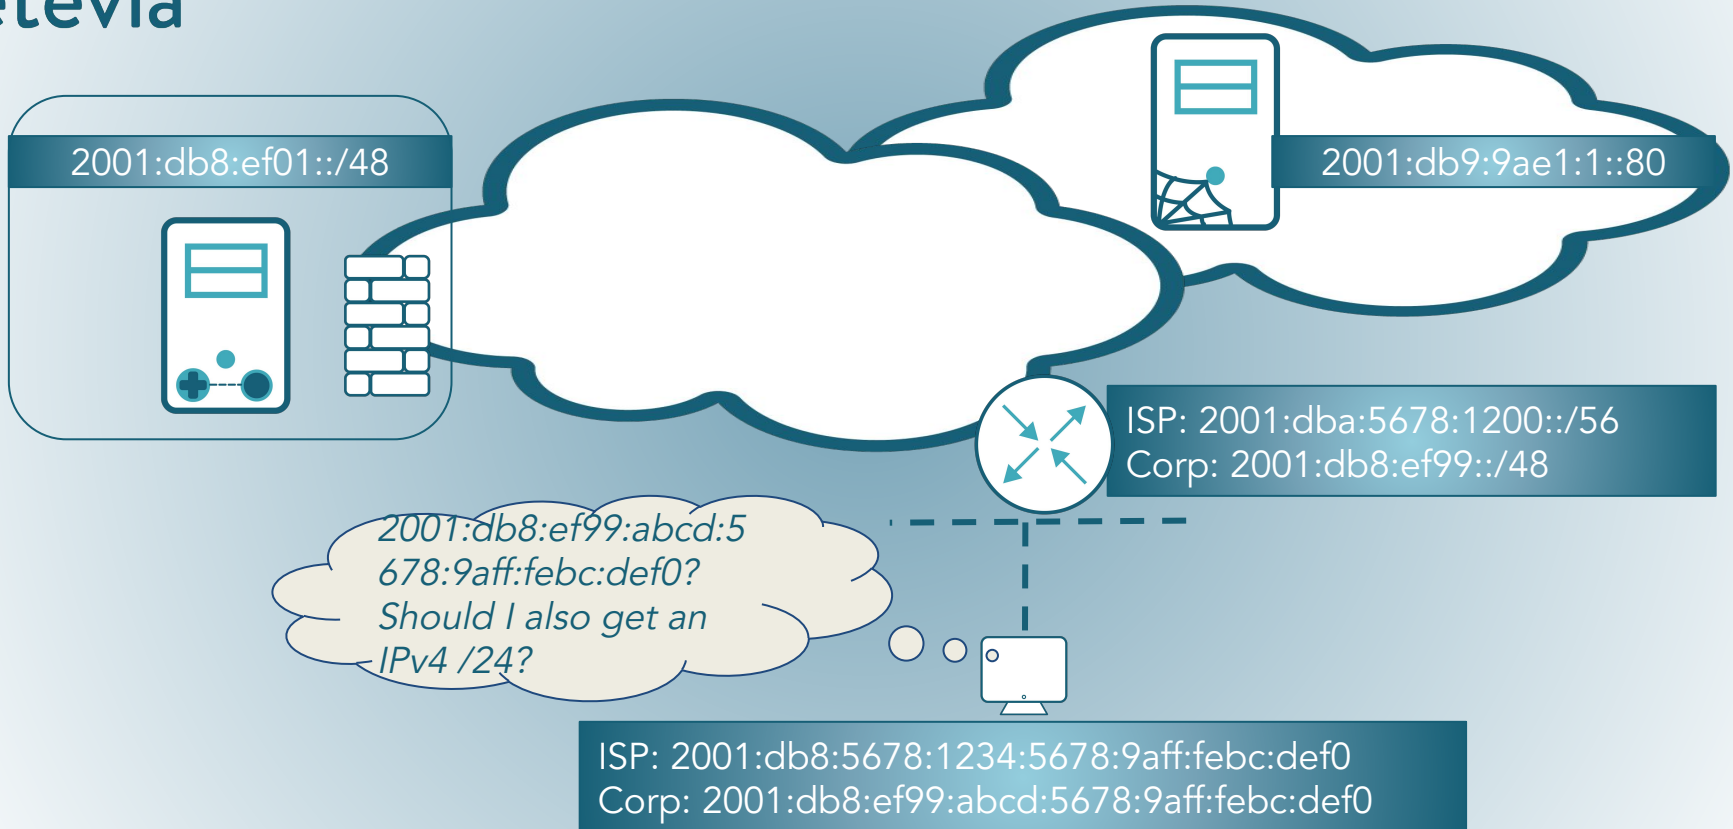
Two Networks



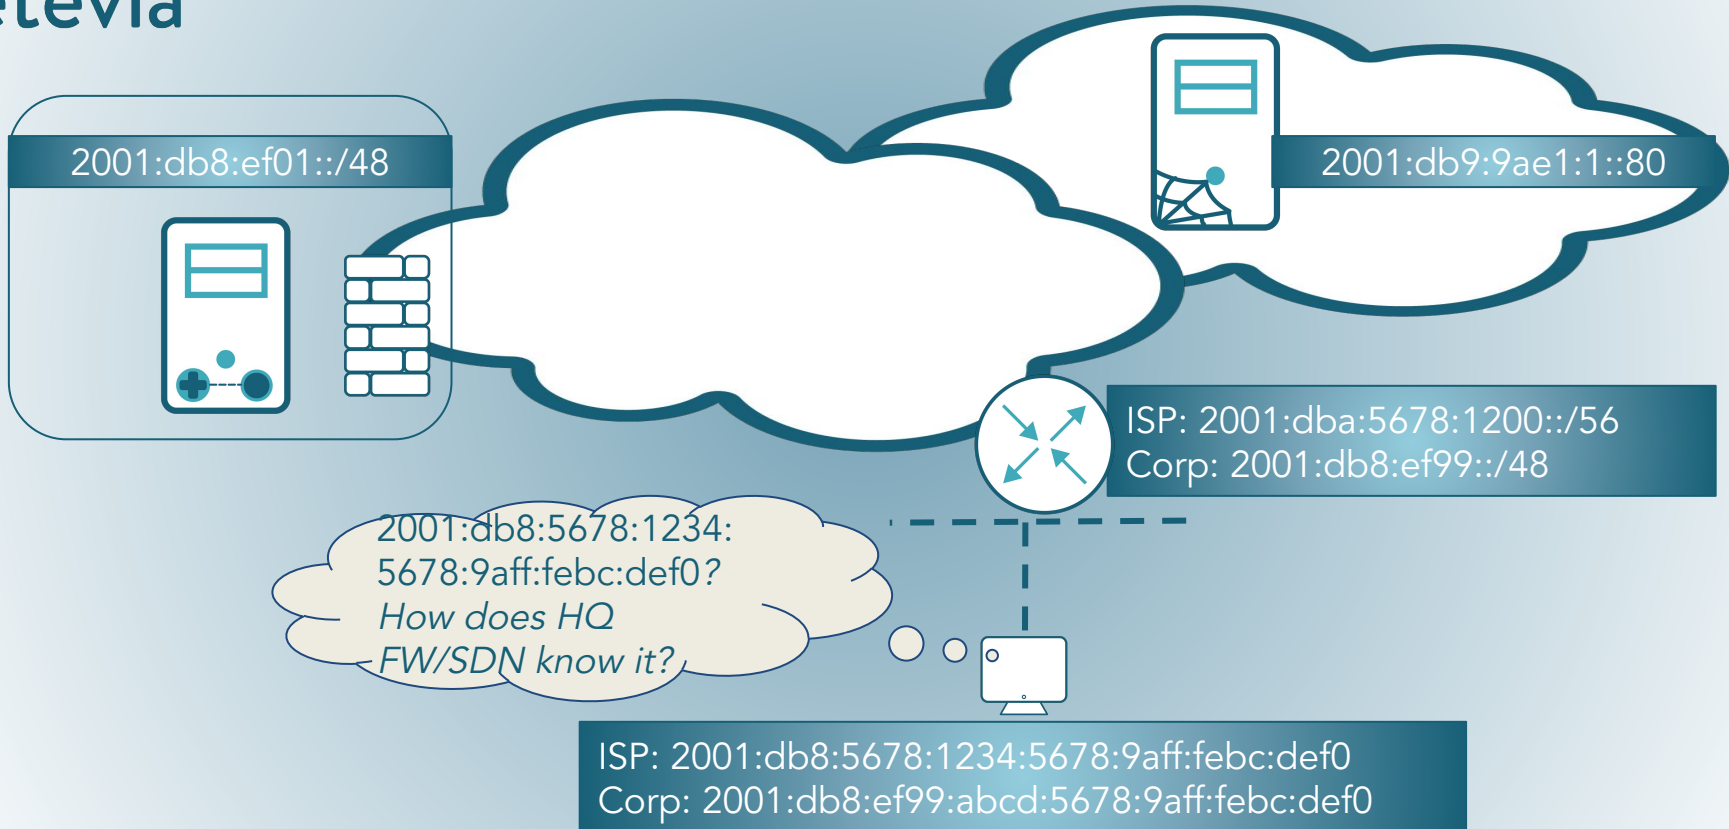
Two Networks



Two Networks



Two Networks



Provisioning Domains

- Provisioning Domain info might include
 - Source address to use in PvD
 - IP addresses of DNS server
 - HTTP proxy (if any)
 - DNS suffixes for the network
 - Default gateway address

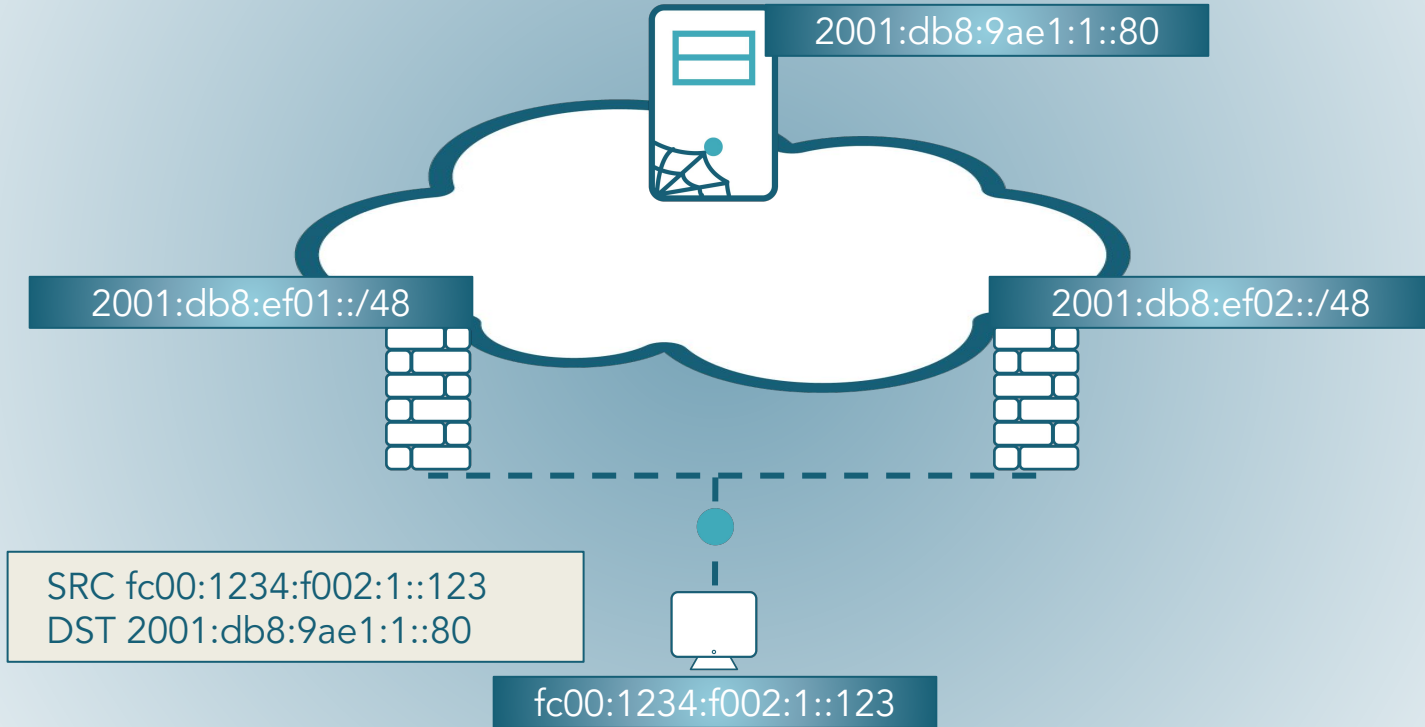
Source Address Selection

1. Avoid unusable destinations
2. Prefer matching scope
3. Avoid deprecated addresses
4. Prefer home address
5. Prefer matching label
6. Prefer higher precedence
7. Prefer native transport
8. Prefer smaller scope
9. Use longest matching prefix
10. Leave order unchanged

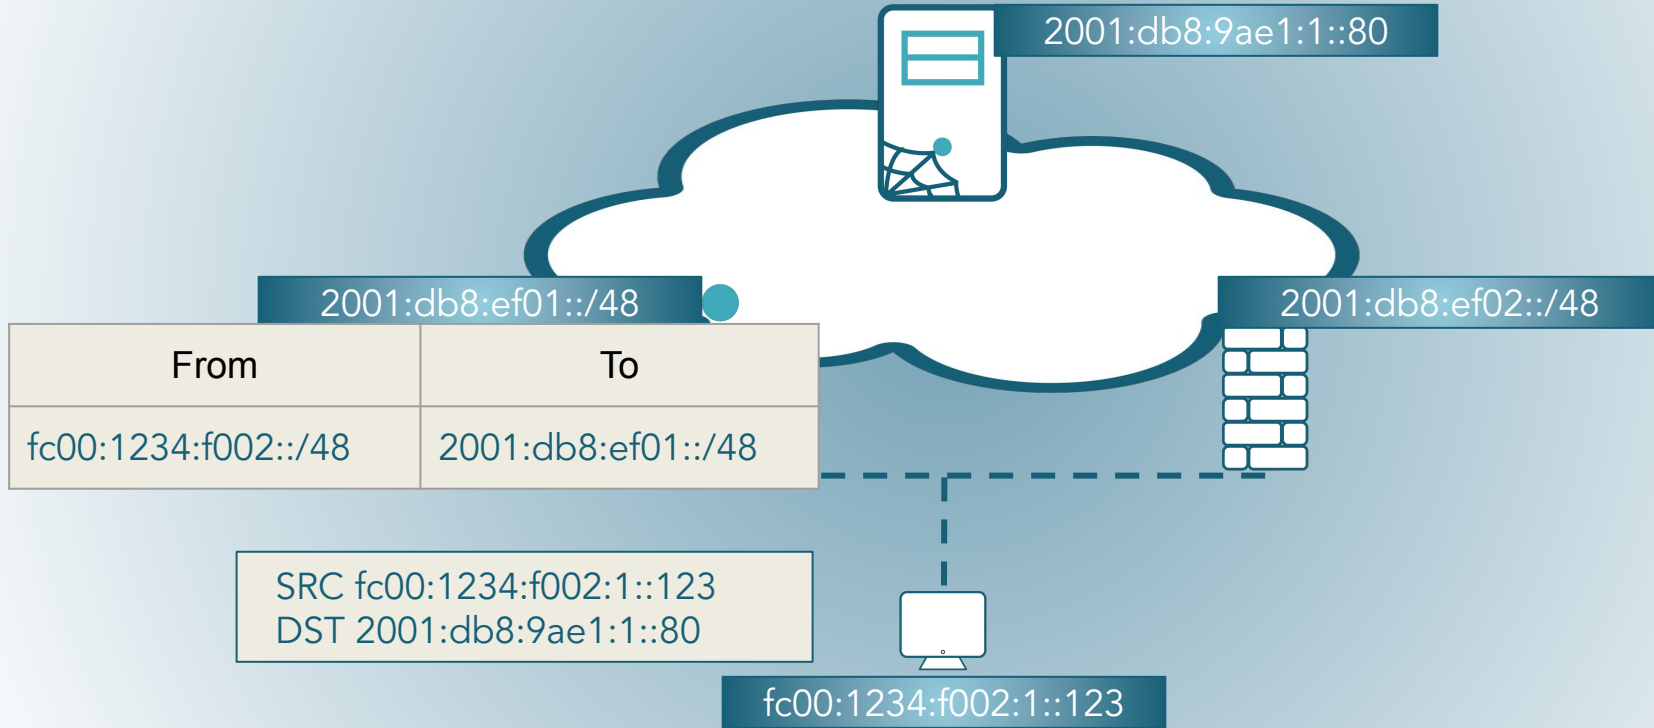
Sorry - this problem isn't solved yet

See <https://tools.ietf.org/html/draft-ietf-intarea-provisioning-domains-02> for leading candidate (identify PvD with a FQDN in the RA)

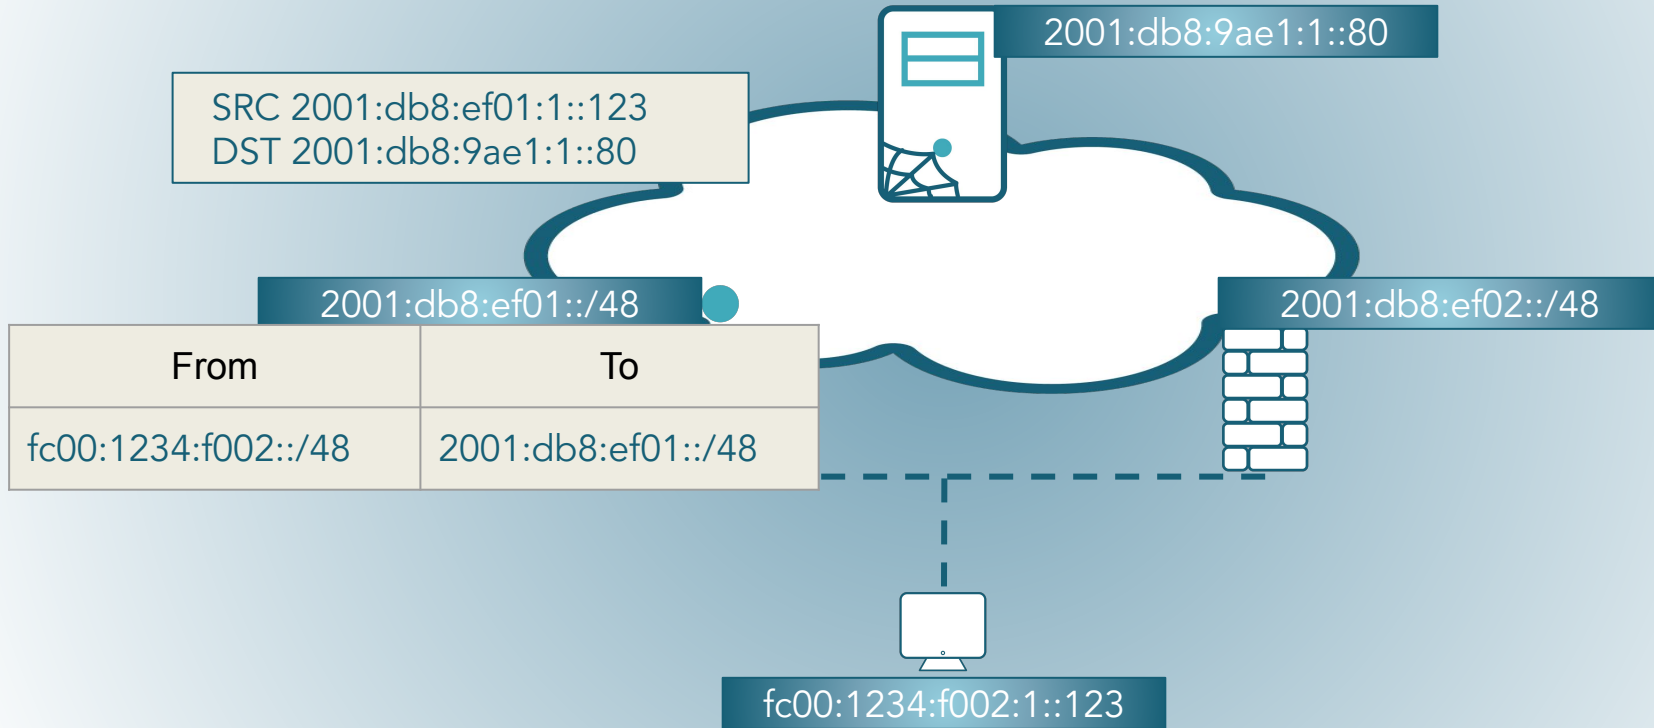
Network Prefix Translation



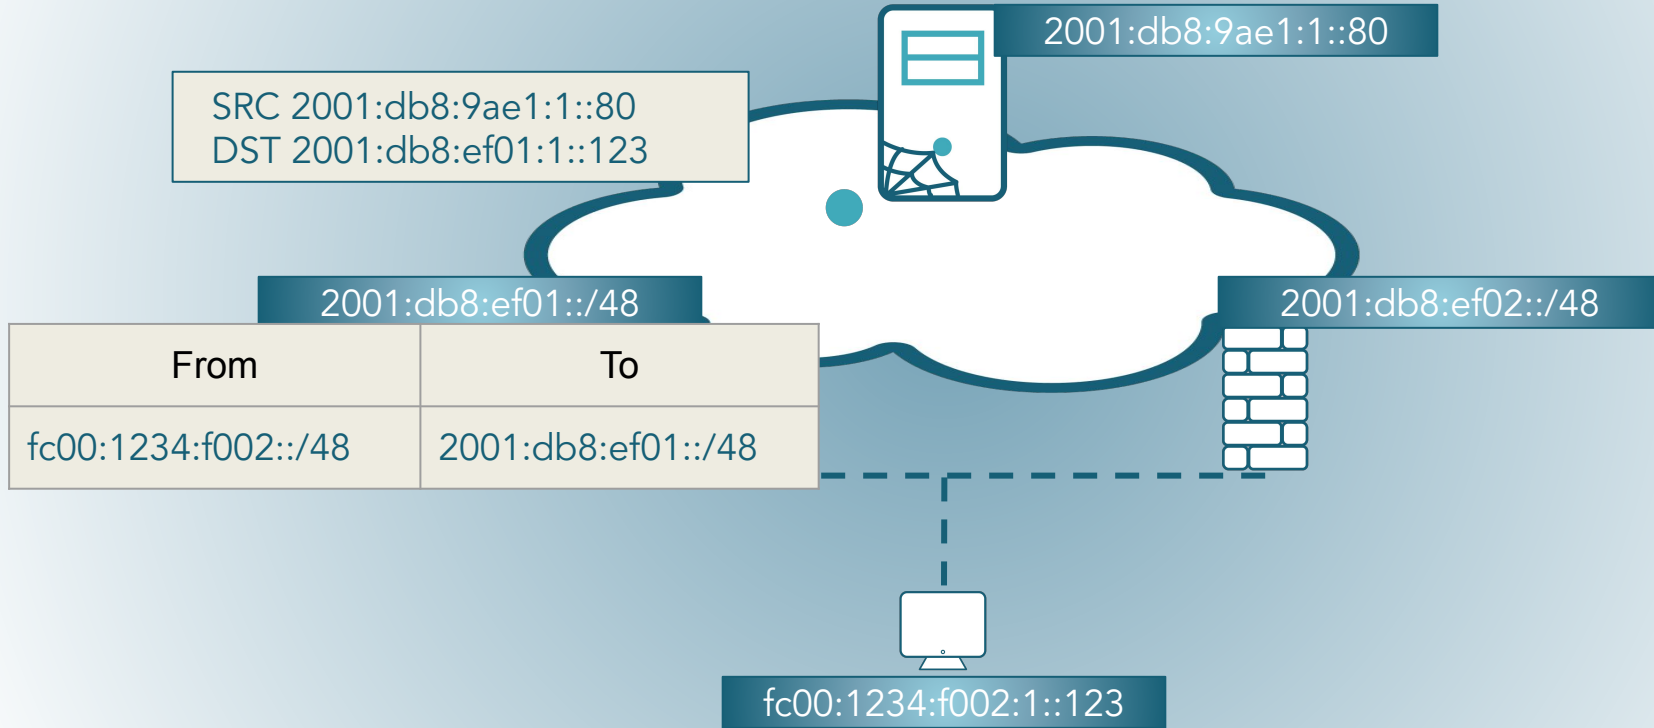
Network Prefix Translation



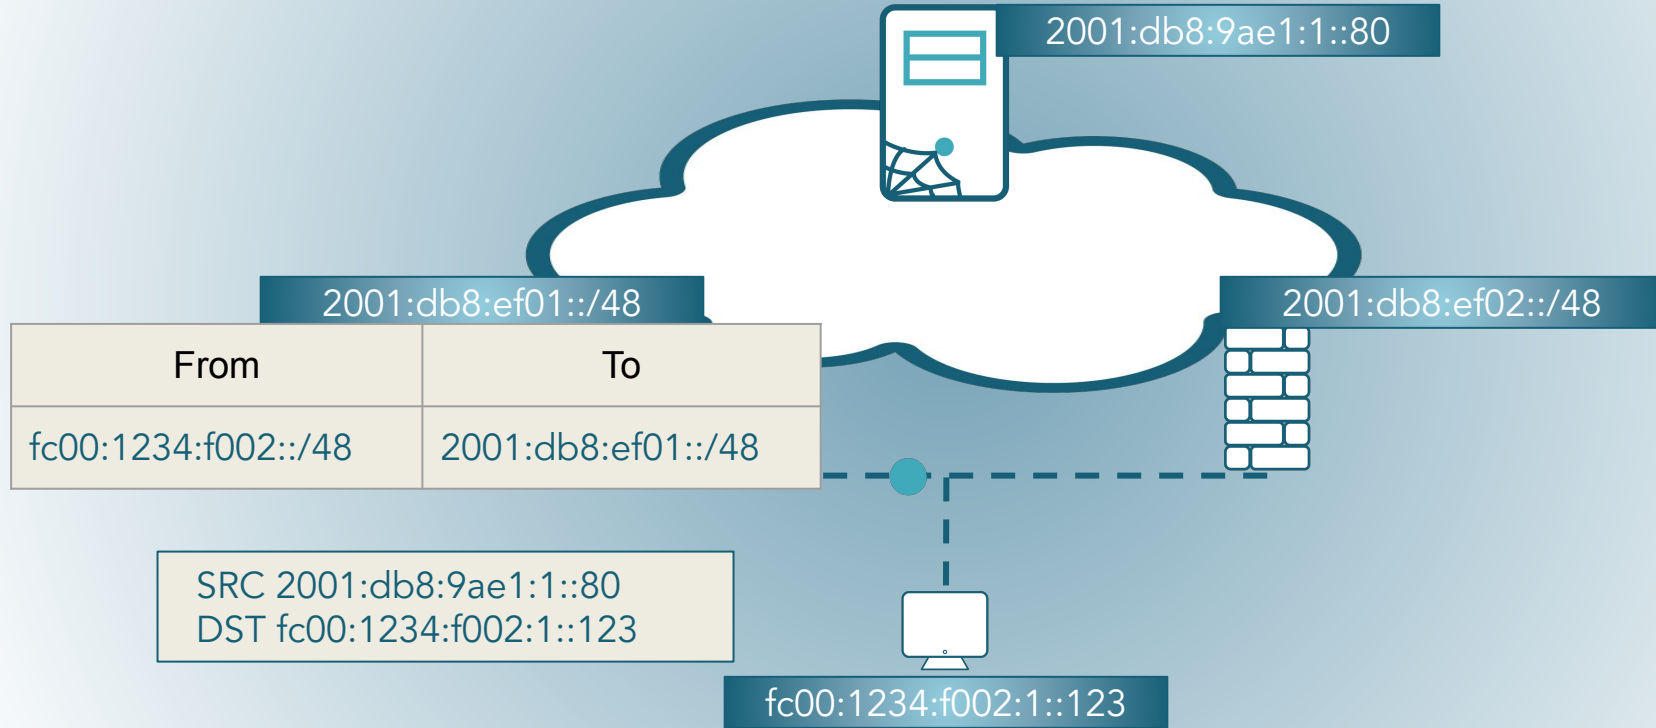
Network Prefix Translation



Network Prefix Translation



Network Prefix Translation



Connecting the Office

What else should I worry about?



Other Obstacles

- Additional considerations for IPv6 deployment (ISPs, devices, web, and what can be done)
- • If not otherwise covered, a summary of reports from Cisco and Microsoft's IPv6-only experiences



retevia.net

Discussion

